


AFRL-SN-RS-TR-1998-37
Final Technical Report
May 1998



HIGH PERFORMANCE SEED BASED OPTICAL COMPUTING

Lucent Technologies

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. A274

19980727 184

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
SENSORS DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-SN-RS-TR-1998-37 has been reviewed and is approved for publication.

APPROVED:



JAMES L. DAVIS
Project Engineer

FOR THE DIRECTOR:



ROBERT G. POLCE, Acting Chief
Rome Operations Office
Sensors Directorate

If your address has changed or if you wish to be removed from the Air Force Research Laboratory Rome Research Site mailing list, or if the addressee is no longer employed by your organization, please notify AFRL/SNDR, 25 Electronic Parkway, Rome, NY 13441-4515. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

HIGH PERFORMANCE SEED BASED OPTICAL COMPUTING

A. L. Lentine, R. L. Morrison, D. J. Reiley
R. A. Novotny, D. B. Buchholz, J. M. Sasian
K. W. Goossen, M. G. Beckman, S. J. Hinterlong

Contractor: Lucent Technologies
Contract Number: F30602-93-C-0166
Effective Date of Contract: 9 August 1993
Contract Expiration Date: 9 August 1996
Program Code Number: 3D10
Short Title of Work: High Performance Seed Based Optical
Computing
Period of Work Covered: Aug 93 - Aug 96

Principal Investigator: Anthony Lentine
Phone: (630) 713-5445
AFRL Project Engineer: James L. Davis
Phone: (315) 330-4276

Approved for public release; distribution unlimited.

This research was supported by the Defense Advanced Research
Projects Agency of the Department of Defense and was monitored
by James L. Davis, AFRL/SNDR, 25 Electronic Parkway, Rome,
NY 13441-4515.

List of Figures	iii
List of Tables	v
Acknowledgments	vi
List of Acronyms	vii
1. SUMMARY	1
2. OBJECTIVES	2
3. INTRODUCTION	3
3.1 Demand For High Bandwidth Telecommunications: Yesterday, today, and tomorrow .	3
3.2 Free space Optical Interconnections	4
3.3 Free Space Optical Interconnection demonstration systems: a Review	6
3.4 The need for long wavelength optical modulators.	10
3.5 The need for increased functionality: Optoelectronic VLSI	10
3.6 VCSELs versus Modulators.	11
3.6.2 Optoelectronic VLSI	11
3.6.3 Modulator based OE-VLSI: Circuit/System Issues.	12
3.6.4 VCSEL based OE-VLSI	12
3.6.5 Conclusion	13
4. TECHNICAL ACCOMPLISHMENTS	15
4.1 Long Wavelength MQW Modulators	15
4.2 Optoelectronic VLSI	19
4.2.2 Capacitance modeling of OE-VLSI.	20
4.2.3 Analysis of Thermal Resistance.	30
4.3 Demonstration System.	49
4.3.1 Introduction	49
4.3.2 Architecture	50
4.3.3 Switching Chip.	54
4.3.4 Switching Chip Mounts	64
4.3.5 Optical System.	68
4.3.6 Diffractive Optics.	83
4.3.7 Mechanical Design.	84
4.3.8 Fiber Bundles.	90
4.3.9 Input Lasers	93
4.3.10 System Control Software.	94
4.3.11 Experimental Results	97
4.3.12 Multichannel optical oscilloscope for sampling optoelectronic circuits	103
4.3.13 Conclusion	109
4.4 System 7.	110
4.4.1 System 7: Goals and Objectives	110
4.4.2 System 7 Switching Chips.	110
4.5 Future Optical Systems	115

4.5.1	Introduction	115
4.5.2	Paraxial Optics of Free-Space Photonic Switching Systems.....	115
4.5.3	Point to point interconnections using diffractive optical systems	127
5.	RECOMMENDATIONS	132
6.	REFERENCES	133
7.	INCLUDED PUBLISHED PAPERS	142

List of Figures

Figure 1	8
Figure 2	9
Figure 3	9
Figure 4	16
Figure 5	17
Figure 6	19
Figure 7	20
Figure 8	20
Figure 9	21
Figure 10	23
Figure 12	25
Figure 11	26
Figure 13	29
Figure 14	32
Figure 15	35
Figure 16	36
Figure 17	38
Figure 18	41
Figure 19	43
Figure 20	45
Figure 21	51
Figure 22	52
Figure 23	53
Figure 24	54
Figure 25	56
Figure 26	56
Figure 27	57
Figure 28	59
Figure 29	60
Figure 31	61
Figure 30	62
Figure 32	65
Figure 33	66
Figure 34	66
Figure 35	66
Figure 36	67
Figure 37	67
Figure 38	67
Figure 39	70
Figure 40	71
Figure 41	72
Figure 42	74
Figure 43	74
Figure 44	75

Figure 45	77
Figure 46	77
Figure 47	78
Figure 48	81
Figure 49	82
Figure 50	86
Figure 51	87
Figure 52	89
Figure 53	94
Figure 54	95
Figure 55	96
Figure 56	97
Figure 57	98
Figure 58	100
Figure 59	102
Figure 60	104
Figure 61	106
Figure 62	107
Figure 63	108
Figure 64	112
Figure 65	112
Figure 66	113
Figure 67	114
Figure 68	117
Figure 69	121
Figure 70	123
Figure 71	123
Figure 72	128
Figure 73	129
Figure 74	130
Figure 74	130
Figure 75	131

List of Tables

Table 1	4
Table 2	4
Table 3	5
Table 4	6
Table 5	7
Table 6	24
Table 7	24
Table 8	33
Table 9	34
Table 10	40
Table 11	46
Table 12	63
Table 13	73
Table 14	76
Table 15	86
Table 16	111
Table 17	122

Acknowledgments

This work was sponsored in part by the Defense Advanced Research Projects Agency (DARPA) under the Ultra Dense Ultra Fast Computing Components Program and by Rome Laboratory (AFMC).

We also greatly acknowledge a host of people at Bell Labs, Lucent Technologies who have contributed to the work presented here. The design, fabrication and material growth of the long wavelength MQW modulators was done by K. W. Goossen, J. E. Cunningham and W. Y. Jan. Process development and fabrication for the Optoelectronic VLSI technology was done by K. W. Goossen, J. A. Walker, L. M. F. Chirovsky, L. A. D'Asaro, B. J. Tseng, S. P. Hui, R. E. Leibenguth, J. E. Cunningham, W. Y. Jan, J. M. Kuo, D. P. Kossives, D. D. Bacon, and D. Dahringer. Modulator characterization was done by A. L. Lentine, R. A. Novotny, K. W. Goossen, and G. Livescu. Circuits were contributed by R. A. Novotny, T. K. Woodward, A. V. Krishnamoorthy and A. L. Lentine. The chip interface flexible circuit was designed by D. B. Buchholz. The architecture was invented by G. W. Richards and T. J. Cloonan. The optical system was designed and tested by J. M. Sasian and D. J. Reiley. The mechanical system was designed by J. M. Sasian, D. J. Reiley, M. G. Beckman and D. Vukobratovich. D. J. Reiley performed most of the assembly and alignment of the system. The diffractive optical elements were designed by R. L. Morrison and fabricated by D. B. Buchholz. The control and testing software was written by R. L. Morrison and A. L. Lentine. Diagnostic tools were contributed by R. L. Morrison, W. Knox, S. G. Johnson and A. L. Lentine and the characterization of the system was performed mostly by A. L. Lentine and R. L. Morrison. The laser drivers and receiver subcircuits were designed, assembled, and tested by R. A. Novotny, T. J. Cloonan, J. Loos and S. J. Hinterlong. The fiber bundle was made by J. M. Sasian. S. J. Hinterlong was responsible for project management.

List of Acronyms

VLSI	Very large scale integrated circuit
CMOS	Complementary metal oxide semiconductor
NMOS	n-type channel metal oxide semiconductor
PMOS	p-type channel metal oxide semiconductor
OE-VLSI	Optoelectronic VLSI
MQW	Multiple quantum wells
SEED	Self electro-optic effect device
S-SEED	Symmetric self electro-optic effect device
FET	Field effect transistor
FET-SEED	Field effect transistor self electro-optic effect device
ECL	Emitter coupled logic
p-ECL	Positive voltage emitter coupled logic
ATM	Asynchronous transfer mode
SONET	Synchronous optical network
ITU	International Telegraph Union
SDH	Synchronous digital hierarchy
IP	Internet Protocol
VPI	Virtual path identifier
VCI	Virtual channel identifier
EGS	Extended generalized shuffle (network)
LAN	Local area Network
NRZ	Non return to zero (data format)
RZ	Return to zero (data format)
MOPA	Master oscillator power amplifier (laser source)
DBR	Distributed Bragg reflector
CCD	Charge coupled device
EDM	Electric discharge machining
SEM	Scanning electron microscope
TEC	Thermal Electric Cooler
GPIB	General purpose interface bus
C_{in}	Input capacitance
C_{amp}	Amplifier capacitance
C_{bump}	Flip-chip bond pad capacitance
C_{diode}	MQW diode capacitance
C_{trace}	CMOS metallic trace capacitance
t_{oxide}	Oxide thickness
t_{bump}	Flip-chip bond pad thickness
t_{metal}	CMOS metallic trace thickness

t_d	Gate delay
t_0	Unloaded gate delay
f_{osc}	Ring oscillator frequency
\AA	angstrom (10^{-10}m)
λ	wavelength
α	absorption coefficient
Δ	exciton linewidth
ϵ_0	permittivity of free space
ϵ_{oxide}	relative dielectric constant of the oxide in the CMOS
R_{TH}	Thermal resistance
R_{total}	Total thermal resistance
R_{spread}	Spreading thermal resistance
R_{bump}	Flip-chip bond pad thermal resistance
R_{SiO_2}	Oxide thermal resistance
k_{solder}	Solder Thermal conductivity
k_{SiO_2}	Oxide thermal conductivity
k_{GaAs}	Gallium Arsenide thermal conductivity
PTAT	Proportional to absolute temperature
E	Elastic Modulus
ΔT	Temperature change
y_o	Object height
y_i	Image height
u_o, u_i	Numerical apertures in the object and image planes
\bar{u}_o, \bar{u}_i	Slopes of the chief rays in the object and image planes
\bar{u}_s	Slope of the marginal ray in the aperture stop.
u_s	Field of view of the objective lens
$f/\#$	f number of the objective lens
A/S	aperture stop

1. SUMMARY

This final report represents the culmination of a three year program at Lucent Technologies sponsored by DARPA and Rome Laboratory, beginning in August of 1993, and ending February, 1997. The objectives of the program, which are more specifically stated on the following page, were to enhance the development of surface normal photonic and optoelectronic components and free space optical subsystems for use by both military and commercial applications. To this end we achieved several significant contributions, which are described in this document.

The silicon based optoelectronic VLSI offered a tremendous breakthrough in so called smart pixel technologies, where one integrates surface normal optical I/O with electronic logic. The integration is based on flip-chip bonding of the III-V optoelectronics devices onto commercial VLSI circuits. The benefit of using commercial silicon VLSI is huge; we can enjoy the yields and performance that accompanies the billions of dollars invested in mature silicon VLSI technologies. In this report, three circuits are described, one for the system demonstrator and two subsequently designed with better performance. The latter circuits have up to 400k FETs, 4000 optical I/O, with per channel and aggregate capacities approaching 1 Gb/s and 1 Tb/s respectively. These circuits are described in sections 4.3.3 and 4.4.2.

The optical system provided imaging between a 2 dimensional fiber array and the single chip. The mechanical design of the system uses a plate-pedestal system, that provides superior robustness compared to the slot-plate systems. This system is mounted in a standard electronic equipment frame. The system contains a single two dimensional fiber array providing fibers for the input signals and read beams and providing fibers for the output beams. The optical system images the inputs from the fiber bundle onto the switching chip, provides optical fan-out of the signals from the fibers to the switching chip, and images the outputs from the chip onto the fiber bundle. A section of the switch using 16 input fibers and 16 output fibers was operated as a 208 Mb/s time multiplexed space switch, which is applicable to ATM switching using the appropriate out-of-band controller. A larger section with 896 input light beams and 256 output beams was operated as a slowly reconfigurable space switch.

2. OBJECTIVES

The goals of the research program as originally intended included the following:

- The development of devices, components and processing technologies that allow the integration of longer wavelength optical and optoelectronic interconnects and I/Os. They include the development of 2D arrays and 3D assembly techniques that are integrable, extensible, and scalable. Device packaging and reliability issues are emphasized.
- The construction of a high performance demonstrator, a packet based interconnection network, that exploits the capabilities of the longer wavelength SEED technology and its required optical hardware.

To that end, the following intermediate objectives were proposed.

Year 1:

- Fabricate and characterize long wavelength quantum well modulators
- Document describing design rules for systems based on SEED technology and free-space optical interconnects

Year 2:

- Demonstrate longer wavelength integrated devices
- Develop optical testing tools (e. g. interferometers, alignment probes, 2-D sampling scope)

Year 3:

- Demonstrate arrays of switching nodes
- Deliver arrays of switching nodes for use in system demonstrator
- Demonstrate representative portion of the 256 x 256 packet based network. In particular, that network should have per channel data rates in excess of 155 Mb/s, an excess of 750 optical I/O, and an aggregate I/O bandwidth exceeding 50 Gb/s.

The program was successful, in that all necessary steps to complete the demonstration were undertaken. The silicon based optoelectronic VLSI technology exceeded all expectations that were present in the monolithic based technology at the time of the contract. The long wavelength devices were fabricated and tested, but improvements in receiver sensitivity, laser power at shorter wavelengths, and sub-system architecture made their adoption into the optoelectronic VLSI technology platform unnecessary. However, we should point out, that Lockheed-Martin has successfully integrated these long wavelength devices, perhaps basing some of their work on some of our modulator results.

3. INTRODUCTION

3.1 Demand For High Bandwidth Telecommunications: Yesterday, today, and tomorrow

The demand for interconnection bandwidth is increasing at a frantic pace, well beyond expectations even a few years ago. The explosive growth of the world wide web and the adoption of the internet by business and non-technical households with users in all age groups seems to be driving this growth. Other examples contributing to the demand for increased bandwidth include high definition television (HDTV), video on demand and medical imaging. This demand necessitates increased performance in several areas in both telecommunication and data communications. Among these are increased transmission bandwidths, increased switching capacity, increased internet router and server capacities.

Fiber optic telecommunications transmission systems began to replace microwave radio systems in the early 1980s. Their advantages were manifold. Perhaps greatest was the increased bandwidth of fiber optics compared to microwave links. The bandwidth of a microwave radio channel was 20-30 MHz, which could support 90 - 150 Mb/s using multilevel encoding schemes such as quadrature amplitude modulation. Also important, was that the fiber channel was not prone to multipath fading that occurs when atmospheric bending of the microwave signals caused destructive interference at the receive antenna. By the mid-1980s, fiber optic systems, now operating at rates beyond 1 Gb/s, had completely replaced microwave radio as the medium of choice for all long haul telecommunications systems.

In North America, fiber optic systems conform to the SONET standards. SONET is the North American Synchronous Optical Network standard for telecommunications transmission using fiber optic cables. It provides a uniform set of protocols for the management of high bandwidth services. It includes a multiplexing structure, optical parameters, service mappings, and operations support for existing and future services. In addition, standardized interfaces allow vendor-independent interconnection of terminal and subsystems. SONET was developed by Committee T1 as a universal transport system. The International Telegraph Union, Telecommunications Standards Committee (ITU-T, formerly CCITT) adopted SONET as the basis for its SDH (Synchronous Digital Hierarchy) transport system. Currently, SONET is the North American subset of the ITU's SDH.

In standardizing interfaces, savings can be accrued in operations expense while simultaneously improving the quality of service delivered to the customer. SONET also provides the underlying infrastructure necessary to support new service offerings.

A summary of SONET and SDH data rates are given in Table 1.

For multimedia and data, two somewhat competing standards have evolved. ATM is becoming a standard for transport of data over SONET links, usually over distances up to and beyond 1 km or so. Asynchronous transfer mode (ATM) is a high-performance, cell-oriented switching and multiplexing technology that utilizes fixed-length packets to carry different types of traffic. ATM is a technology which will enable carriers to capitalize on a number of revenue opportunities through multiple ATM classes of services, high-speed local area network (LAN) interconnection, voice, video, and future multimedia applications in business markets in the short

US	Europe	Bit Rate (total)
STS-1	--	51.84 Mb/s
STS-3	STM-1	155.52 Mb/s
STS-12	STM-4	622.08 Mb/s
STS-24	STM-8	1244.16 Mb/s
STS-48	STM-16	2488.32 Mb/s
STS-192	STM-64	9953.28 Mb/s

Table 1. SONET and SDH designations and data rates [1].

term and in community and residential markets in the longer term.

On the other hand, Ethernet has emerged as the leading standard for local area networks for connections between computers within buildings. Other LAN types include Token Ring, Fast Ethernet, Fiber Distributed Data Interface (FDDI), Asynchronous Transfer Mode (ATM) and LocalTalk. Ethernet is popular because it strikes a good balance between speed, cost and ease of installation. These strong points, combined with wide acceptance in the computer marketplace and the ability to support virtually all popular network protocols, make Ethernet an ideal networking technology for most computer users today. Ethernet links operate at 10 Mb/s, 100 Mb/s and, most recently, 1 Gb/s.

3.2 Free space Optical Interconnections

The increased bandwidth of transport products and demonstrations has greatly exceeded the progress in switching. The main reason for this is that switching is very computationally intensive and that progress in electrical I/O bandwidth both from PC-boards and from chips has been slow compared to the increase fiber optic capacities. Electrical interconnects, among other things, are limited by loss in the conductor; this frequency dependent loss limits the achievable bit-rate. The loss can only be reduced by making the lines thicker and wider, and this tends to reduce the number of lines per unit length, thus electrical interconnections will have a maximum bandwidth per unit area [2]. Using equalization can improve the overall bandwidth, but consumes quite a bit of power [3]. In rough numbers, it becomes quite difficult to have IC's with electrical interconnection bandwidths approaching a terabit per second.

Limitations of Electrical Interconnects

- Limited number of connections to circuit boards (partitioning)
- Limited bandwidth of connections (length dependent)
- Crosstalk
- Power dissipation from line drivers and terminations (>20 mW/line)
- Signal and clock skew
- High cost

Table 2. Limitations of electrical interconnections

Optical interconnections offer a chance to circumvent these loss issues. It is likely that a gradual progression of optical interconnects will evolve within large switching and computing systems. We have seen the use of optical data links, operating from 50 Mb/s on up to a Gb/s or more, within digital systems. Several standards now exist, most notably fibre-channel with interconnection bandwidths of approximately 1 Gb/s. These optical data links are suitable for connections between printed circuit boards, not yet for connections between chips.

Features of Optics for Interconnections

- High frequency of Optics
 - No frequency dependent loss or crosstalk
 - Intrinsically very high bandwidth medium
- Short Wavelength
 - Essentially no distance dependent loss or degradation
 - Possibility of high aspect ratio interconnects
 - Possibility of large numbers of interconnections without clock skew in two dimensional arrays
 - Possibility of Global Interconnect topologies
- Large Photon Energy
 - Electrical Isolation
 - Immunity to electromagnetic interference
 - Fundamentally lower communication energy

Table 3. Features of Optics for interconnections [4]

Beyond this, several companies and government laboratories are now pursuing parallel optical data links. These parallel links contain arrays of receivers and lasers, allowing up to 10 Gb/s or so bandwidth between printed circuit boards.

These near-term approaches to optical interconnects will likely grow within the next few years. However, there is a demand, even today, for systems whose performance will not be met by the use of parallel links on PC cards. The use of optical interconnections normal to the surface of optoelectronic chips offers the potential for further increased bandwidths. While the potential to solve a host of design problems exists (see Table 4[3.3]), several very important pieces need further development before insertion into systems. These are:

- An optoelectronic VLSI technology (OE-VLSI), integrating optical detectors and lasers or modulators onto electronic circuitry. This step is most important, for optical interconnects to compete with electronic interconnects, you must gain in interconnect features and performance without losing in processing power or performance
- Low cost, high reliability optomechanical packaging. Reliability here is very important, for

Specific Design Problems Solved by Array Optical Interconnection Technology

- On-chip
 - electro-static discharge protection on inputs
 - simultaneous switching noise
 - area of off-chip drivers and pads
 - power dissipation of off chip drivers
- Chip-to board
 - pin inductance on chips and chip carriers
 - limited number of pins
 - large number of power and ground pins
- On-board and board-to-backplane
 - impedance matching
 - line termination
 - electrical isolation
 - crosstalk between lines
 - bandwidth limits of lines

Table 4. Table 3.3. Specific Design Problems Solved by Array Optical Interconnection Technology [4]

electrical system designers are not willing to take risks in reliability.

- Low cost fiber-optic interfaces - if the system in question connects to other equipment via fibers.

The progress made under the direction of this contract in all three of these areas is indeed impressive.

3.3 Free Space Optical Interconnection demonstration systems: a Review

The history of using free space optical interconnects for optical computing is a long one. While Bell Labs announced the "first optical computer" in 1990 [5], decades of work pre-dated this [2]. However, the invention and manufacture of arrays of self electro-optic effect devices (SEEDs)[112,70] were a major contribution of the Bell Labs effort, because it was the first system to be built with semiconductor devices that would have high speed and later prove to be easily integrated with electronic circuitry.

The work at Bell Labs was split into four locations, the work of our group in Naperville IL concentrated on applying free space optical interconnections to telecommunications switching

equipment. The nature of telecommunications switching was ideal for free space photonics; it required the concentration of several hundred to thousand high speed optical inputs, routing the data from these inputs to their correct outputs, and required the output data in optical form. We never really did solve the total problem optically, because the electronic problem of data routing was too great.

The system described in this report is the sixth system designed and constructed in this arena. The progression has included improvements in architecture, devices, optics, and mechanics. As a general rule, subsequent systems have operated at higher data rates, with fewer components, and with greater mechanical stability as the systems have progressed. In the paragraphs that follow, we will briefly describe the systems.

System	Year	Fabric Size	Array Size	Stages	Bit Rate	Switching Element	Optics	Mechanics	Ref
1	1988	2 x 2	2 x 1	2	10kb/s	S-SEED	Catalog	Microbench	[7]
2	1989	32 x 32	4 x 8	4	10kb/s	S-SEED	Catalog	Pedestal	[12]
3	1990	64 x 64	8 x 16	3	100kb/s	S-SEED	Catalog	V-groove	[118]
4	1991	16 x 32 32wide	32 x 32	6	100kb/s	S-SEED	Custom	Slot-plate	[117]
5	1993	32 x 16	4 x 4	5	155 Mb/s	FET-SEED	Custom	Slot-Plate	[104]
6	1996	16 x 16 x 16	same	1	208 Mb/s	OE-VLSI	Custom	Plate-Pedestal	[14]

Table 5. History of free space switching system demonstrations at Bell Laboratories from 1988-present.

The first system, if you can call it that, showed the implementation of a pair of 2 input 1 output switching nodes using 4 Symmetric SEEDs on two arrays [7]. The S-SEEDs [13] could implement either set-reset latches or logic gates depending on how they were operated. The experiment was constructed using commercial microbench mechanics and off the shelf optics.

System 2 [12] consisted four stages of S-SEEDs implementing part of an extended generalized shuffle network (EGS) [8]. The system used off the shelf optics mounted on pedestals on a flat plate. Control of the network was static, index cards in paths of the clock beams were used to set up routing paths, but the goal was to use a spatial light modulator in the long run.

System 2 had several weaknesses, but most notable was the difficulty aligning the optics. Thus, in the system 3 [118], the optical components were mounted in v-grooves for easy centering of the lenses. Plossl relay lenses allowed the field of view to be expanded from 4 x 8 devices to 8 x 16 devices. The devices still used S-SEED arrays, control was still provided by blocking sets of beams.

A couple of interesting experiments were performed after the initial demonstration. In one experiment, a section of the array was operated at 1 Mb/s. In a second experiment, a Logic SEED [15], consisting of 10 quantum well diodes, was used in place of two device arrays to implement the switching function in a single stage.




Figure 1. Several of the components from System 2 (on the left) and System 3 (on the right.)

System 3, which followed the Bell Labs computing system by only 6 months was a major hit at the time with a live demonstration at Supercom '90. In its configuration there, it performed the function of a lens, merely inverting (and repeating) the data from input to output, but later we did actually perform switching with the network.

In these first three systems, we agonizingly tried to reduce the optical losses in the systems. As a result, the combination and separation of the input data signals, read power beams and output beams was performed by an intricate arrangement of patterned mirrors and relay pairs. [15]. Additionally, the shuffle interconnect required for the networks, was performed by an equivalent crossover network [9]. While in theory, these networks were "lossless", in practice, both optical losses and aberrations proved to be major problems limiting the field of view (size of the switching array) that one could achieve. As a result, our fourth system [117], provided a drastic simplification in the optical system. The system used a custom beam splitter with a 50/50 mirror to provide the beam combination/separation and a 1×3 binary phase grating with blocked orders to implement a banyan interconnection network. The system used a custom objective lens imaging 8000 light beams onto the 2 mm field of view of the chip. It still used S-SEEDs as the switching element, but this time the control of the network was accomplished by adjusting the bias on columns of devices via a personal computer. The mechanical design of the system consisted of a number of slots on a plate, (an extension of the v-grooves). This design was later used by a number of research groups worldwide [see e.g. 11].

In spite of the improvements in optics and mechanics, the systems still had a problem in that their overall data rate was low. Indeed the S-SEED was not capable of operation much beyond 50 Mb/s, because the required optical energy was too high. Also, independent control of 1000s of channels was difficult, because, while the S-SEED was a memory element, it was difficult to bring in additional beams just to hold the state of a memory. Hence, it had become clear, several years earlier actually, that electronics was needed both to improve receiver sensitivity and to enhance functionality.

System 5 [104] implemented a 5 stage network using FET-SEEDs as the active elements. FET-SEEDs consisted of the monolithic integration of multiple quantum well optical modulators




Figure 2. System 4 (on the left) and System 5 (on the right.)

and detectors with GaAs field effect transistors. It enabled memory elements for control and improved the required optical energy from a few picojoules in the S-SEEDs to about 0.1 picojoule in simple FET-SEED receivers. The optical system was further improved to use pupil division, which made larger spots at the receivers, but allowed once again, lossless beam combination.

The data rate goal of System 5 was 155 Mb/s. System 5 never did achieve this very well, primarily to the difficulty of designing FETs with uniform threshold voltages. In our first attempt at it, the bias lines were made too narrow, and voltage drops on the lines caused non-uniform receiver thresholds. This limited our data rate to ~50 Mb/s. In our second attempt, we achieved 150 Mb/s operation in a couple of channels, but a threshold voltage too far below our design voltage caused drastically reduced noise margin in our buffered fet logic gates; as a result, not all channels functioned correctly.

System 6 [14], the subject of most of this report, had improved mechanics, optics, and switching devices. The mechanics was based on a plate-pedestal system that provided improved stability compared to the slot-plate systems of systems 4-5. The optics, while conventional, had a




Figure 3. System 6 (on the left) and components for System 7 (on the right.)

field of view greater than 7 mm, enabling us to image thousands of spots onto devices with larger windows and greater functionality than any of our previous efforts. A custom fiber bundle was

designed and manufactured that had both inputs and outputs on the same bundle. The architecture required only a single stage, eliminating the complication of the multiple devices in our previous systems. It was very stable, operating at 155 Mb/s for more than 8 months without realignment.

System 7, which will be described briefly in this report, has a drastically reduced size compared to System 6 and was to have operated at 622 Mb/s. Data on the switching arrays showed error free operation at 780 Mb/s, and open eyes beyond 900 Mb/s, so it is likely that it would have met the objective. As of this writing System 7 is incomplete.

3.4 The need for long wavelength optical modulators

At the time of the beginning of the contract, most quantum well modulators were designed and fabricated for operation at 850 nm using GaAs/AlGaAs materials. The fact that the GaAs material in the well was binary made epitaxially growing material easier than a ternary or quaternary material. Operation at 850 nm required 10 nm wells, and this width gave a reasonable compromise between the oscillator strength (height of the absorption peak versus wavelength) and the shift with applied field. That is, wider wells shift faster with field, but the absorption peak is less pronounced. Single mode semiconductor lasers were readily available at 850 nm, although their optical powers were < 30 mW at that time. The types of optoelectronic systems that we were building required laser powers in excess of 1 Watt, and such high power lasers were not available at 850 nm at that time. High power Nd: YAG lasers at 1.06 micrometers were readily available, and high power monolithic power amplified lasers (MOPAs) at 980 nm were becoming available. Hence the desire to move to one of these wavelengths.

However, several developments contributed to the decision to pursue optoelectronic VLSI technology at 850 nm. These are 1) Potential availability of high power lasers, both solid state using Cr:LiSaF, and semiconductor (MOPAs), at 850 nm, 2) A sub-system architecture, comprised of a single switching chip, that not only required lower laser powers per input but allowed the laser power to be supplied via multiple fibers. 3) Improved receiver sensitivity using silicon VLSI versus the monolithic FET-SEED circuits.

3.5 The need for increased functionality: Optoelectronic VLSI

Prior to the contract, Bell Laboratories had developed a smart pixel technology platform consisting of the monolithic integration of GaAs FETs and multiple quantum well detectors and modulators[70]. Circuits were designed, fabricated and tested consisting of a 4 x 4 array of 2 input 1 output switching nodes. Each chip had 96 quantum well diodes and 400 FETs. Operation of individual circuits was demonstrated above 400 Mb/s [16].

The monolithic technology was limited in several ways. First, it was very difficult to control the threshold voltages of the FETs, particularly since we were not piggybacking on an established GaAs technology. Thus, from wafer to wafer, the performance of the circuits varied. This variation forced us into extremely simple receiver designs, that achieved switching energies of roughly 100 fJ at 155 Mb/s. Second, the yield of the individual FETs was poor in comparison to silicon VLSI; this yield, although still greater than 99%, kept us from fabricating functional circuits with 1000s or 10000s of FETs. Third, the circuits were designed in buffered FET logic, which led to very high static power dissipations, also keeping the circuit complexity small. Direct coupled FET logic (DCFL) would have been better, but a 'true' enhancement mode device was

unavailable. Fourth, there was a performance penalty of ~25% associated with the monolithic integration. Submicron CMOS had better performance than our GaAs circuits, and the performance gap was likely to increase because of the limited resources that we could allocate to improving the GaAs process.

So it was clear to us (and to others in the field) that a technology integrated with silicon would be superior. We began an effort in monolithic integration of III-V materials on silicon, but there were several drawbacks to that approach. These include the possibility of damaging the submicron FETs with the temperatures required by the epitaxial growth and the need to complete the submicron metallization after growth. Because of these, and other fabrication related issues, we pursued the hybrid integration of silicon VLSI with quantum well modulators and detectors.

The resultant technology, which will be explained in more detail in section 4.2, has the following attributes compared to the monolithic GaAs based technology.

- 1) Because the CMOS is commercial, circuits with nearly 100% yield with millions of FETs can be designed.
- 2) Because much of the circuitry is 'static' and CMOS has low static dissipation, power dissipation per functional unit is an order of magnitude less than a comparable circuit in the monolithic FET-SEED technology
- 3) Because the threshold variations are bounded and the local uniformity in threshold is good, receivers have been designed with an order of magnitude better sensitivity compared to the monolithic technology.
- 4) The raw performance of the silicon circuits is faster. Gate delays for 0.35 micrometer CMOS are better than 50 psec, whereas our GaAs had gate delays of greater than 120 psec.
- 5) We have achieved 4000 optical I/O and others have achieved 60k optical I/O, owing to the fact that the yields of the optoelectronic circuits are not limited by electronics.
- 6) We have designed more functional 'pixels', implementing complete 16 x 16 switches rather than simple 2 x 1 switches.

3.6 VCSELs versus Modulators

As we stated earlier, during the next decade the demand for bandwidth between circuit packs and chips will continue to increase. Parallel optical data links consisting of linear arrays of VCSEL based transmitters connected by fiber ribbons to arrays of receivers [17] is a leading approach to providing bandwidth beyond today's single optical data link solutions. However, this approach falls short beyond perhaps a hundred fibers or so, necessitating the need for two-dimensional arrays of optoelectronic devices. The VCSEL technology appears to be a natural for these two dimensional arrays because of the simple fact that it emits light normal to the surface.

3.6.2 Optoelectronic VLSI

A particularly useful platform for two dimensional optical interconnects is Optoelectronic VLSI (OE-VLSI). OE-VLSI consists of the hybrid integration of (at least) thousands of optical

sources or modulators, optical detectors, and silicon VLSI circuitry. While a VCSEL based OE-VLSI platform is currently desired by many working in the field, a MQW modulator based platform [18,19,20,21,22] is currently *commercially available* [20] with Gb/s I/O speeds[21], up to 65k optical I/O [22] and potential aggregate I/O bandwidths exceeding a terabit per second [23]. VCSEL based OE-VLSI [24,25], has recently achieved the integration of 256 VCSELs on a silicon driver, with only one VCSEL operating concurrently, and no integrated detectors [25]. The technological challenges in manufacturing a usable VCSEL based OE-VLSI platform are formidable, yet, the desire by many for this technology over the modulator based technology remains. Quantitative analyses have been performed that compare the overall dissipation of VCSEL and modulator based systems [26,27,28]. This section will examine, qualitatively, the issues of building systems with both technologies, assuming of course, that a VCSEL based OE-VLSI platform will one day exist.

3.6.3 Modulator based OE-VLSI: Circuit/System Issues

The detriment of having to supply an array of bias beams to read their state really depends on the system. In a recent MQW OE-VLSI system demonstration [14], the bias beams were provided by fibers, increasing the total number of fibers by only 6% compared to what a VCSEL based system would require. Using multiple fibers for the read beams eliminated the need for a single high power laser source.

Detailed design trade-offs for MQW modulators [14] show the potential for devices with a 2:1 contrast ratio over 17 nm or 60C at 3V, using active bias stabilization[30]. Higher contrast ratios or greater tolerance to wavelength and temperature variations can be achieved by operating at greater detunings from the exciton peak [31], using thicker devices operating at higher voltages. By using multiple stacked MQW modulators [32] that operate at low voltages but high fields, one can design low voltage MQW modulators with improved contrast and tolerances at the expense of increased capacitance. Alternatively, one can use high voltage drivers[33] without an increase in capacitance, however, either of these approaches will increase the power dissipation at a given bit-rate.

Modulator based systems offer at least three options that VCSEL based systems cannot provide. First, modulator based systems have the ability to do re-timing at the modulator output for multistage systems, when combined with a dynamic receiver at the input. The receiver will require lower average powers when operated with pulsed read beams [34,35]. For simple functions, this avoids the need to electrically distribute the clock; effectively, one can do retiming of multiple chips with sub-picosecond resolution. Second, modulator based OE-VLSI output devices can be programmed to be either detectors or modulators. This allows bidirectional data flow on a single channel [36,37]. Lastly, using a resonant MQW modulator, logic levels can be represented as equal amplitude phase states[38], allowing the possibility for reconfigurable holographic interconnections[39,40].

3.6.4 VCSEL based OE-VLSI

It's been said that "VCSELs are to optoelectronics as CMOS is to electronics[41]." While we have much experience in modulator based systems, several issues aren't well understood in VCSEL based systems.

Obviously, VCSELs do not require an array of bias beams, thus simplifying the optical systems. However, if the VCSELs are operated in the regimes that produce multiple transverse modes, the optical system must operate at higher numerical apertures than for a single mode system, and this precludes the use of single mode fibers or waveguides. This perhaps *complicates* certain optical systems compared to modulator based ones. Also, modulators can be aligned simply to single mode fibers by making them larger than the fibers in 'retroreflector systems' [42].

Circular VCSELs have arbitrary polarization, thus tend to preclude the use of polarization components, and there will be some polarization aberration in any optical system that will lend itself to excess noise[43]. Asymmetrizing the cavity promises to reduce this polarization switching [44,45,46,47,48], but during turn-on and turn-off, one might experience some polarization as well as modal variations. A truly single mode VCSEL is the best good solution, but it requires very small VCSELs and requires reduction of various loss mechanisms [49].

The VCSEL promises high contrast, but it comes at the expense of delay variations dependent on the bit pattern; these delay variations ultimately limit the bit-rate [50]. These delay variations depend on the on and off state currents relative to the threshold current. The highest speed VCSEL experiments have their off-state bias above threshold, which limits the contrast ratio and potentially gives large off state power variations. For the highest speed operation, the on-state bias must be many times threshold, and VCSELs often support multiple transverse modes at currents beyond few times the threshold current. On the plus side, zero bias VCSEL links ($\sim 250 \mu\text{A}$ threshold) have been recently reported with 1 Gb/s data rate and only a 1 dB penalty versus the same link with the VCSEL biased [51].

Thermal variations across an array of VCSELs can cause threshold shifts which cause variations in optical power and turn on delay [52,53,54,55,56]. Thermal crosstalk has been a huge issue in monolithic arrays of VCSELs[55], it could be worse in flip-chip bonded arrays, because the thermal conductivity between devices should be greater. However, if sub-milliwatt operation is feasible, thermal crosstalk should not be a limiting factor. Wavelength variations, due to thermal effects or material growth, are also likely to be in the few nm range [56], limiting the use of diffractive optical fan-out in VCSEL based systems.

3.6.5 Conclusion

In this section, we've outlined some of the issues that we've faced in the MQW based OE-VLSI systems that we've built, and described what we feel are unresolved issues in VCSEL OE-VLSI based systems that will be built in the future. In the modulator based systems that we've developed, the receiver design and performance dominates the complexity of the OE-VLSI chip. It dissipates the most power, it is most susceptible to crosstalk, it has the most design options, it (in some cases) limits the bit-rate and it provides the largest opportunity for improvement in the bit-rate - tolerance - dissipation trade-off. The logic circuitry, depending on its complexity, might also have high power dissipation and limit the bit-rate. In MQW OE-VLSI systems, the performance of the MQW modulator hasn't limited the performance of the system. VCSEL based OE-VLSI systems would also still be limited by receiver and circuit issues.

A still larger issue on insertion of the technology into 'real systems' is the lack of a suitable packaging infrastructure that is palatable to electronic system designers. In ten years of progress in optical interconnects, there are only a handful of packaging schemes applicable to 2D

optoelectronic technology, and none seem viable for insertion into, for example, a desktop PC.

Most impediments to the adaptation of widespread use of parallel optical interconnects and many of the improvements that are required, both in packaging and on chip, are independent of the choice of output device. It is important that system and circuit level research supplement device level research to move this technology to the point where it can be applied on a massive scale.

4. TECHNICAL ACCOMPLISHMENTS

4.1 Long Wavelength MQW Modulators

Because of its ability to provide high power with high spectral and spatial quality, the Nd:YAG laser, operating at 1.064 nm, has been considered as a possible light source for free space photonic systems. MQW modulators at these longer wavelengths can be made, in principle, using either GaAs or InP substrates. Since the GaAs substrate is a better candidate for large scale integration than InP, much attention has been focused on InGaAs/GaAs MQW's for this application [57,58,59,60,61]. However, since the MQW must be at least 1 μm thick for useful surface-normal modulation, strain relief is bound to occur in this system. This relaxation of the lattice results in dislocations which propagate upward resulting in a striated surface [60]. This surface roughness results in unwanted diffraction of the light beams. In addition, the defects make the integration of GaAs transistors problematic.

We have shown that the InGaAs/GaAsP material system may be used to grow undefected MQW's on GaAs for modulators at 1.064 μm [62,63]. This is because the addition of phosphorous to the barrier results in negative strain in the barrier balancing the positive strain in the InGaAs well, allowing in principle any number of periods to be grown without having any net strain buildup. Our devices are optically smooth and the sharpness of X-ray scattering spectra indicate that no relaxation of the lattice occurs. This material system can provide MQW's with band gaps anywhere from 870 nm to 1.064 μm . We have also shown that modulation-saturation is high for modulators using the MQWs, allowing operation as high as tens of kilowatts per square centimeter [64].

It has typically been found by ourselves and others that the absorption coefficient of longer wavelength MQW's is reduced, by about a factor of two for excitons at 1064 nm, compared to GaAs/AlGaAs MQW's with excitons at 850 nm. Since modulators rely on changes in absorption coefficient, the corresponding modulation is also reduced by about a factor of two. Before this work, it was not clear whether the reduction in absorption coefficient was intrinsic to the long-wavelength material systems or due to broadening of the excitons. Here we show that the latter is the case. We measured absorption coefficient (α) and linewidth (Δ) of samples with excitons from 850 to 1064 nm, and found the product of α (when normalized to the total well thickness) and Δ to be roughly constant. Thus, the integrated absorption is constant, and by reducing the linewidth of long-wavelength modulators, their performance should become equal with that of GaAs/AlGaAs devices. Of course, this work does not show whether the broadening is an intrinsic effect (e.g., due to alloy scattering, or possibly, as discussed below, ultrashort exciton lifetimes).

Five long-wavelength p-i(MQW)-n modulators were grown using the InGaAs/GaAsP system on n-type GaAs substrates using Gas-source molecular-beam epitaxy. The mole fraction of In was varied from 0.11 to 0.24 to produce excitons ranging from 920 nm to 1064 nm. Correspondingly, the P content in the barrier was adjusted from 0.6 to 0.75 to maintain a strain balanced condition, resulting in defect free samples, which was checked by X-ray diffraction. The substrates were not rotated during growth, resulting in each sample covering a range of wavelengths. In Fig. 4, the samples from the same substrate are identified by an ellipse about their spectra. Each sample had 50 wells which were 95 \AA wide for the samples with excitons below 990 nm, 90 \AA for the sample with excitons between 1000 and 1040 nm and 85 \AA wide for the 1064 nm sample. The

barrier width was 60 Å for all the samples except the 1064 nm sample which had 65 Å barriers. Atop each sample, a 5000 Å thick p-type GaAs layer was grown, 200 μm x 200 μm mesas were etched on each sample, and gold contacts were made to the p layer. The backside of each sample was polished and antireflection coatings were applied to both the front and back surfaces.

For comparison, p-i(MQW)-n samples were fabricated using the GaAs/AlGaAs system. A




Figure 4. Absorption spectra of various MQW modulator samples at 0 volts bias normalized to the entire intrinsic width. Spectra grouped by ellipses are from the same unrotated wafer

sample with a 1 μm thick bulk GaAs intrinsic layer was fabricated. A sample with a conventional GaAs/Al_{0.3}Ga_{0.7}As MQW with 86 Å wells and 44 Å barriers (60 periods) was made. Finally, a shallow GaAs/Al_{0.02}Ga_{0.98}As MQW sample, [65] with 100 Å wells and 40 Å barriers (56 periods) was made. All these samples had their substrates removed and were antireflection coated for transmission measurements.

The absorption spectra-(normalized to the entire intrinsic width of the sample) of all the samples is shown in Fig. 4. For the long-wavelength samples, the transmission was normalized to unity at a detuning of 120 meV from the exciton peak. For the GaAs/AlGaAs samples, slight residual Fabry-Perot fringes made establishing the absorption baseline somewhat non quantitative, and can be observed in Fig. 4. As can be seen in Fig. 4, the absorption coefficient is greatly reduced for the long wavelength samples compared to the deep GaAs/AlGaAs MQW. This is partly due to the fact that the barriers are wider for the long wavelength samples, adding optically inert material to the MQW. In Fig. 5, the absorption coefficient, normalized only to the total well thickness in each sample (α) is plotted (open squares). This parameter should more fairly compare samples as if their barriers widths were equal. Note, however, that reducing the barriers of the strain balanced long wavelength samples would require higher strain (more P) in the barriers in order to maintain the strain-balanced condition.

We observe in Fig. 5 that α for the long wavelength samples is still greatly reduced compared to the deep GaAs/AlGaAs MQW. However, we also plot in Fig. 5 the linewidths (HWHM, Δ) of the samples' excitons. The linewidth of the long wavelength samples is much larger than that of the deep GaAs/AlGaAs MQW. When the product of absorption coefficient (α) and linewidth (Δ) is plotted, we see that it remains roughly constant for all the samples. There is a slight

rise in the product as wave length increases, which may be due to increased electron-hole overlap due to the higher barriers, although since the shallow wells show strong exciton with nearly zero confinement, this explanation may be incorrect. Since this product represents the spectrally integrated absorption, this value should remain unchanged if the linewidth of the long wavelength samples were somehow reduced. Therefore we indicate here that if the linewidths of long wavelength MQW excitons can be reduced to that of GaAs/AlGaAs MQW's, their corresponding absorption coefficients should increase to become comparable also, since the integrated absorption should be constant. Note that the constancy of the absorption-linewidth product does not extend to other material systems. In [66], $\alpha = 0.94 \text{ 1}/\mu\text{m}$ (normalized to total well width) and Δ (HWHM) = 4.5 meV, giving an $\alpha\Delta$ product of 4.23 meV/ μm , less than half of what we find for our samples. (Thus it would appear that InP/InGaAs MQW's have reached the limit of their performance.)

The exciton broadening may be inhomogeneous in nature, i.e., caused by nonuniformities




Figure 5. Absorption coefficient (normalized to the total well width - open squares), linewidth (x's) and their product for the samples of Fig. 4. The integrated absorptions of the excitons of all the samples are roughly the same.

in the sample. The nonuniformities can be in the form of, for example, alloy fluctuations, or as another example, interface roughness. If the latter is the dominant cause, then reduced linewidth may be achieved by improved crystal growth. However, the former is usually thermodynamic in nature, and so probably cannot be controlled greatly by growth conditions. If this is the case, it may not be possible to ever achieve the linewidth of GaAs/AlGaAs MQW's, which have nonalloy wells.

The exciton broadening may be homogeneous in nature, i.e., caused by ultrashort exciton lifetimes. This is of some concern due to fact that we recently observed large saturation intensities in InGaAs/GaAsP MQW modulators at 1064 nm [64]. This would tend to indicate that faster exciton ionization (escape from the well) is occurring in these samples, an astonishing possibility given the large band offsets in the system (about 1 eV). However, the effect of lifetime on broadening is probably only small. For shallow GaAs/AlGaAs MQW's, escape from the well has been shown to occur with a single phonon collision [67], which has a time constant of 300 femtoseconds [68]. Since the lifetime induced broadening is given by $\Delta_{\text{life}} = \hbar/(2\tau)$, the linewidth of

shallow quantum wells should be at least 1.1 meV broader than that of deep wells. GaAs/AlGaAs wells, which have much longer escape times. We see in Fig. 5 that the linewidth of shallow wells is indeed about 2.5 meV broader than deep wells. Since escape from the well probably occurs faster in shallow quantum wells than any other sample, it is unlikely then that more than 1-2 meV of exciton broadening can be attributed to lifetime effects for the long-wavelength samples. In conclusion, we have shown that the product of absorption coefficient (when normalized to the total well thickness) and exciton linewidth is roughly constant for MQW modulators with exciton wavelengths from 850 to 1064 nm in the GaAs/AlGaAs and InGaAs/GaAsP material systems. Therefore, reducing the linewidth of the long-wavelength samples should result in an increase of their absorption coefficient to that observed for GaAs/AlGaAs. If the broadening is not due to intrinsic effects (e.g., alloy fluctuations or ultrashort exciton life times), improvement of crystal growth could result in longer wavelength-modulators having performance at least approaching that of 850 nm modulators.

4.2 Optoelectronic VLSI

A particularly useful platform for two dimensional optical interconnects is Optoelectronic VLSI (OE-VLSI). OE-VLSI consists of the hybrid integration of (at least) thousands of optical sources or modulators, optical detectors, and silicon VLSI circuitry. As we've stated above, the advantages of integrating with silicon VLSI are numerous. Because the CMOS is commercial, circuits with nearly 100% yields with millions of FETs can be designed. Because much of the circuitry is 'static' and CMOS has low static dissipation, power dissipation per functional unit is an order of magnitude less than a comparable circuit in the monolithic FET-SEED technology. Because the threshold variations are bounded and the local uniformity in threshold is good, receivers have been designed with an order of magnitude better sensitivity compared to the monolithic technology. The raw performance of the silicon circuits is faster. Gate delays for 0.35 micrometer CMOS are better than 50 psec, whereas our GaAs had gate delays of greater than 120 ps. We have achieved 4000 optical I/O and others have achieved 60k optical I/O, owing to the fact that the yields of the optoelectronic circuits are not limited by electronics. We have designed more functional 'pixels', implementing complete 16 x 16 switches rather than simple 2 x 1 switches.

The technology borrows heavily from focal plane array technology, that has been prevalent and steadily improving for several years. In some ways, our key contribution in process development has been the adaptation of substrate removal, that allowed us to use 850 nm MQW diodes. However, we were the first to apply it to optical interconnections, and so the application of the technology has really advanced since Bell Labs' involvement.

The detailed process for the optoelectronic VLSI [72] is shown in Figs. 6-8. An array of diodes is fabricated in a quantum well p-i-n diode wafer grown by molecular beam epitaxy. Barrier metals and solder are deposited on the array bond pads of the diode array and the silicon chip. The two chips are aligned in a commercial bonder and epoxy is applied to the areas between the chips. The epoxy holds the chips in place, and it protects the top (bottom after bonding) surface of the diodes from the substrate removal. Thus, the process uses a thermal compression bond, rather than reflow. The entire GaAs substrate is removed using a wet chemical etch, and the diodes are anti-reflection coated on what is now the top surface. The epoxy over the silicon circuit is selectively removed (not shown) to allow access to electrical wire bond pads.




Figure 6. GaAs/AlGaAs multiple quantum well diode and silicon circuit after solder deposition and before bonding [72].




Figure 7. Optical Modulator array (top) bonded to silicon VLSI circuitry (bottom) before substrate removal [72].




Figure 8. Finished device after substrate removal [72].

Fig. 9 shows a more detailed view of the optoelectronic chip after bonding. In our current devices, the bond pads are designed in third level metal, and circuitry is actually placed underneath the pads, similar to that shown in the figure. In the circuit designed for the system demonstrator, only a two layer process was used and no circuitry was placed beneath the pads. This reduced the density compared to what was achieved in later circuits.

4.2.2 Capacitance modeling of OE-VLSI

Introduction

The three basic building blocks of a smart pixel circuit are an optical-to-electrical receiver, an electronic logic block, and an electrical-to-optical transmitter. It is advantageous to reduce the propagation delay through each of these blocks in order to avoid additional register logic to operate at high data rates. It is also desirable to ensure a constant temperature over the surface of the chip. This is because the characteristics of the electrical-to-optical transmitters are usually sensitive to changes in temperature.

To reduce the propagation delay through the smart pixel receiver block, it is essential to reduce the front end capacitance (C_{in}). C_{in} has three main components: the photodiode active area, the amplifier input, and the stray interconnect capacitance. The FET-SEED technology min-

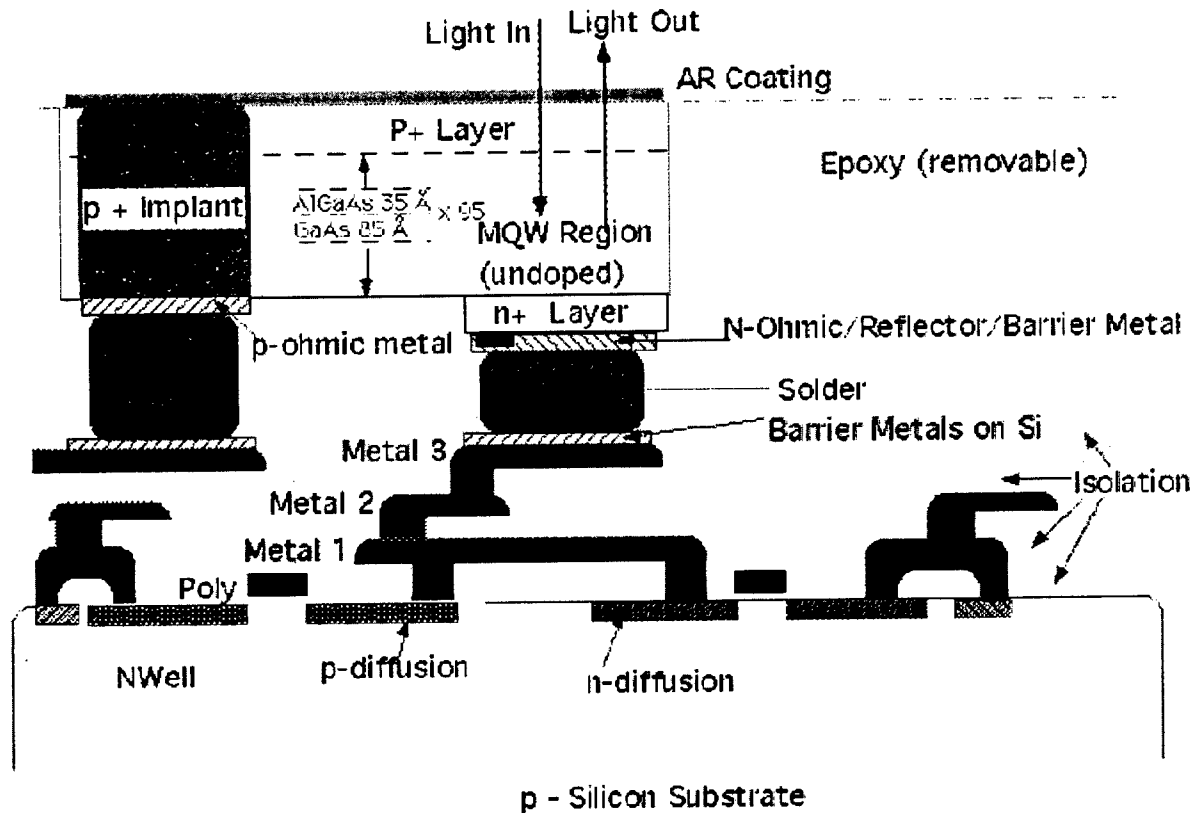


Figure 9. More detailed view of the Optoelectronic VLSI chip after bonding [109].

imized stray interconnect capacitance through the monolithic integration of photodetectors, modulators and electronic circuitry [69,70]. However, the monolithic FET-SEED technology was limited to only medium scale integration (MSI) smart pixel arrays due to yield and power dissipation issues. Hybrid integration of very large scale integration (VLSI) Si CMOS electronic circuitry with photodetectors, modulators, or emitters is an attractive approach in obtaining VLSI smart pixels in the near term.

One method of attaching III-V devices to Si CMOS is through the use of a flip-chip solder bump process and back illuminating the photodiode [71]. A technique, briefly described in Section 1.6, was devised where GaAs MQW diode detectors/modulators were first flip-chip-bonded onto Si CMOS, and the GaAs substrate etched away, to allow operation at 850nm [72]. This process resulted in MQW diode 'islands' on the surface of the silicon VLSI chip.

This section examines the issues of capacitance and thermal conduction of the solder bumped MQW diode islands within the hybrid FET-SEED technology. Section 4.2 investigates the parasitic capacitance as a function of solder bump geometry. Section 4.3 examines the effect of thermal conduction from the MQW diode to Si substrate. Section 4.4 introduces a novel temperature compensated modulator bias circuit that may be used to cancel effects of global temperature changes in optoelectronic chip temperature. Section 4.5 is a summary.

Figure 10(a) depicts a cross-sectional SEM view of the flip-chip hybrid. To begin analyzing this solder bump-bonded diode structure, a simplified model was first constructed and is

shown in Figure 10(b) along with the equivalent circuit. The hybrid process has gone through many revisions, but one common feature throughout is that the bonding pads be equally sized squares spaced one pad width apart. A typical pad size is 15mm [73]. Changes have been made in the size of the diode active area relative to the pad size. When this capacitance study was first performed, the diode active area was defined by a mesa etch and was equal to the size of the pad with an additional 2mm overhang on one side for the ohmic contact [74]. The capacitive loading effect of a diode, less the 15 μm x 15 μm bonding pad contribution, was calculated to be 32fF using data taken from diode loaded ring oscillators [75]. This agreed favorably with the original simulation result of 34fF. Since this time, the geometry of the MQW diode 'island' has been changed in order to improve performance and yield. For instance, the N-mesa is now larger with a size of approximately $(d+3) \times (d+4)$, where d is the dimension of the solder bonding pad in microns. The overall size of the GaAs island has grown too, resulting in additional chip overlap and fringing capacitance. This new structure was analyzed using simple formula estimations and three-dimensional finite element analysis, and measured using diode loaded ring oscillator circuits. Section 4.2.1 presents an updated set of formulae to allow quick estimation of the capacitance. Section 4.2.2 presents the results of a three-dimensional finite element analysis simulation. Finally, Section 4.2.3 presents results of diode loaded ring oscillators fabricated in a 0.9 μm CMOS technology.

Estimation of Capacitance

A formula set was derived to allow quick estimation of the effect of solder bump geometry on the total capacitance for flip-chip solder-bumped technology. The total capacitance was broken down into components depicted in Figure 10(b). The component labeled C_{amp} is the lumped amplifier input capacitance which includes the amplifier Miller effect. C_{diode} is the overlap and fringing capacitance of the photodiode active area. The total stray interconnect capacitance (C_s) is defined here as the sum of the capacitance contributed by the wiring trace (C_{trace}) from pad to amplifier input, the solder bump bonding pad (C_{pad}), the solder bump (C_{bump}), and the GaAs chip (C_{chip}) that forms the bridge between the diode and far pad. Since all of these elements are in parallel, the total front end capacitance, (C_{in}) is the sum:

$$C_{\text{in}} = C_{\text{amp}} + C_{\text{diode}} + C_{\text{trace}} + C_{\text{pad}} + C_{\text{bump}} + C_{\text{chip}} \quad (4.1)$$

where,

$$C_{\text{amp}} = 35\text{fF} \quad (4.2)$$

$$C_{\text{diode}} = (114\text{aF})[1.15(d+2)(d+4) + 4.2d + 16] \quad (4.3)$$

$$C_{\text{trace}} = <2\text{fF} \quad (4.4)$$

$$C_{\text{pad}} = \epsilon_0 \epsilon_{\text{oxide}} [(1.15d^{2/t_{\text{oxide}}}) + (5.6(t_{\text{metal}}/t_{\text{oxide}})^{0.222} + 4.12t_{\text{oxide}}(t_{\text{metal}}/t_{\text{oxide}})^{0.728})] \quad (4.5)$$

$$C_{\text{bump}} + C_{\text{pad}} = \epsilon_0 \epsilon_{\text{oxide}} [(1.15d^{2/t_{\text{oxide}}})] \quad (4.6)$$




Figure 10. Cross-sectional view of solder-bumped MQW diode on silicon substrate
 (a) SEM photo, (b) diagram depicting the equivalent circuit. (not drawn to scale)

$$+ (5.6(t_{\text{bump}}/t_{\text{oxide}})^{0.222} + 4.12t_{\text{oxide}} (t_{\text{bump}}/t_{\text{oxide}})^{0.728}]$$

$$C_{\text{chip}} = \epsilon_0 \epsilon_{\text{epoxy}} [(1.15(d-3)(d+6)/t_{\text{bump}}) + (2.8(d-3)(1/t_{\text{bump}})^{0.222})] (4.7)$$

The formulae used to estimate each element are listed as equations 4.2 to 4.7. The lumped capacitance for C_{amp} and C_{trace} are listed as typical values, to show their relative contribution to the total capacitance. The actual values would depend on the specific circuit design and layout. The formulae for C_{diode} , C_{bump} , C_{pad} , and C_{chip} were derived from the empirical formula set of Sakuri and Tamaru for a plate over a ground plane [76]. The C_{diode} estimates the contribution of the MQW diode active area and includes fringing for the outer 3 sides and 2 corners. The value C_{pad} estimates the bonding pad contribution without solder, while $C_{\text{bump}} + C_{\text{pad}}$ includes the

solder bump. The difference between these two calculations determines the C bump component. The C chip represents the contribution of the GaAs chip between the pads which includes overlap and fringing capacitance for 2 sides. The permittivity constant ϵ_o is $8.85\text{fF}/\mu\text{m}$, and ϵ_{epoxy} is the relative permittivity of the epoxy. The epoxy is used to hold the GaAs and Si chips together during the substrate removal process. A typical value for ϵ_{epoxy} is 2.9. The value t_{metal} is the thickness of the pad metal, and t_{oxide} and ϵ_{oxide} are the thickness and relative permittivity of the oxide and any dielectrics that exist between the pad metal and silicon substrate. The total height of the bump interconnect metal is t_{bump} . Note that these estimates are based on the “as drawn” device structure, and do not include variations introduced during the fabrication process.

PARAMETER	0.8 μm MOSIS	0.9 μm vendor
ϵ_{oxide}	3.9	4.4
t_{oxide}	1.1 μm	0.9 μm
t_{metal}	1 μm	0.5 μm

Table 6. Typical values for the constants shown in Equations 4.1-7

Table 6 shows typical parameter values used in Equations 4.2-7 for first layer metal pads for a 0.9 μm CMOS vendor and the 0.8 μm MOSIS HP CMOS26B process [77]. Figure 11 plots the estimated C in (less the fixed amplifier and trace contributions) vs. pad size for both the 0.9 μm and 0.8 μm process. The results indicate that the bonding pad was the dominate contributor to C_s . Typical solder bump heights from 3-10 μm were found to induce little change on C_s , caused by the trade-off between the C_{bump} and C_{chip} components.

The data presented in Figure 11 assumes the bonding pads were fabricated in first level metal. It is interesting to examine the capacitance of pads formed in higher level metals. Since these metals are further from the substrate, the pad-to-substrate capacitance decreases, however, any underlying metals would serve to increase the capacitance. The effect of the underlying metals is evident in Table 7, which shows the calculated C_{pad} values for a 15 μm x 15 μm pad fabricated in first, second, and third level metal for the 0.8 μm process [78].

15 μm x 15 μm C pad over			
PAD METAL	substrate	1st level	2nd level
1st level	9.6fF		
2nd Level	5.5fF	10.6fF	
3rd level	4.2fF	5.9fF	10.8fF

Table 7. Calculated capacitance values for a 15mmx15mm pad for 0.8mm process

Simulation of Capacitance

To check the accuracy of the approximations, a 3-D Laplace/Poisson solver [79] was used

to calculate the total input capacitance vs. pad size for a MQW diode bumped to the first layer metal on a Si wafer. The results are shown in Figure 12. The estimated values which are reshowed in Figure 12 appear to have overestimated the total capacitance by approximately 10%-15%. This was expected, since the formula set from which the estimations were derived assuming a stand-alone plate over a continuous dielectric and ground plane. The actual structure here was broken down into several stand-alone pieces, not all of which reside over a homogeneous dielectric plane. Some of the side and corner fringing components within the equation set were reduced to account for the overlap between all the pieces, however, the estimate still produces a conservative value. A conservative estimated value is desirable when designing a circuit, to allow some margin for variations in parameters.

The simulations also indicated that the currently used common cathode configuration depicted in Figure 10 resulted in a slightly higher capacitance than a common anode configuration ($\sim 8\%$ higher for a $d=15\mu\text{m}$, $h=3\mu\text{m}$ structure). This was found to be especially true for small bump heights, and was the result of the C_{chip} component adding more significantly to the input capacitance in the common cathode case.

The effect of the epoxy on capacitance was also simulated. A $d=15\mu\text{m}$ pad structure with $h=3\mu\text{m}$ bump height without epoxy was found to have 5.5% less total capacitance than the same structure surrounded by epoxy. As the bump height is increased, the ratio of the capacitance of structures with epoxy to those without epoxy becomes even greater. This effect is due to the increase of the chip and bump capacitive components, caused by the dielectric effect of the epoxy.




Figure 12. Plot of simulated input capacitance for four bond pad sizes of $5\mu\text{m}$, $10\mu\text{m}$, $15\mu\text{m}$, and $20\mu\text{m}$ assuming a $5\mu\text{m}$ bump height. For comparison, the estimated capacitance is also shown as a dotted line.

Experimental Results

To verify the above simulations, 19-stage CMOS ring oscillators were fabricated with and without first level metal solder bonding pad loads. A ring oscillator consists of an odd number of inverters connected in a chain, with the output of the last inverter fed back to the input of the first. When powered up, the feedback causes the circuit to begin to oscillate. By measuring the fre-

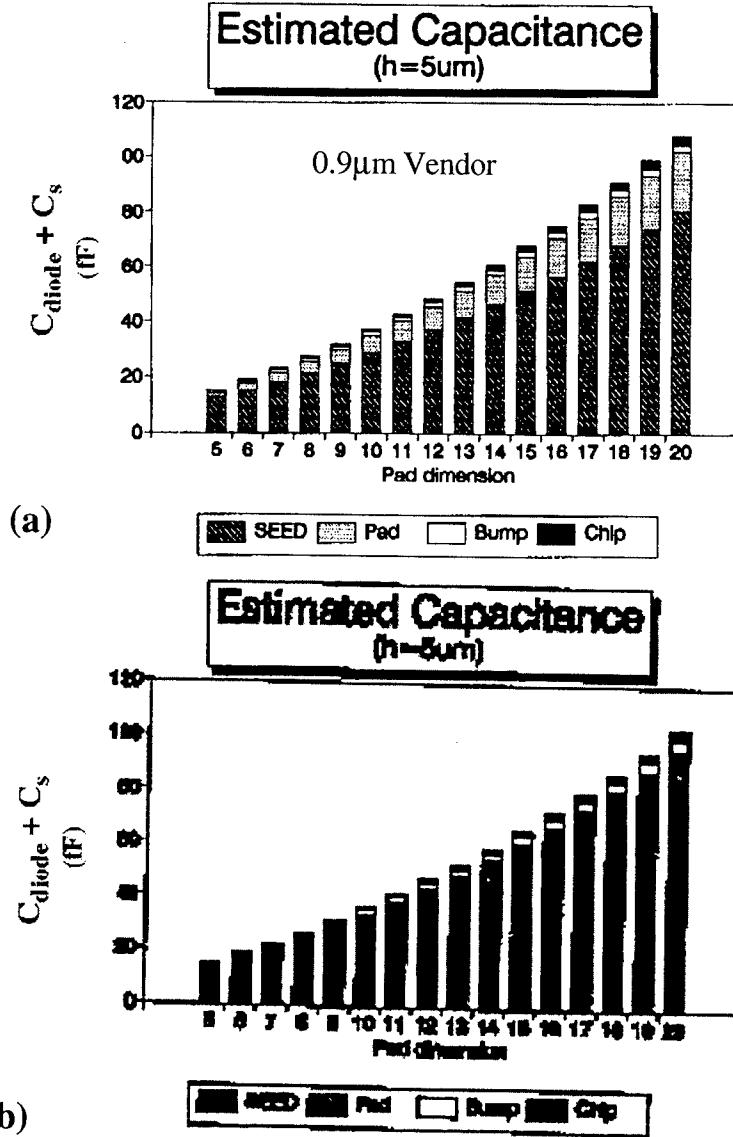


Figure 11. Plot of estimated input capacitance $C_{diode} + C_s$ as a function of bond pad size showing relative contributions of each component (a) for 0.9mm CMOS vendor and (b) for 0.8mm MOSIS HP CMOS26B process. The 0.9mm CMOS process has a higher pad contribution. Note the dominant contribution of the diode active area to the total capacitance in both cases.

quency of oscillation f_{osc} , the average gate delay (t_d) may be calculated using the relation $t_d = 1/(2nf_{osc})$, where n is the number of inverters in the chain. When additional capacitance is added between each inverter, the gate delay increases, and may be approximated using the empirical relation:[80]

$$t_d = t_0 + k C / w_n \quad (4.8)$$

where t_0 is the intrinsic fixed internal delay of the gate, w_n is the width of the N-FET in microns,

C is the load capacitance in fFs, and k is the empirical factor in ps/fF. This formula assumes that the P- and N-FETs are sized within each inverter for equal rise and fall times. The constants t_0 and k may be derived using extracted device parameters along with an analog simulation tool such as SPICE or by measuring the gate delay difference between two ring oscillators with two different known capacitive loads.

A bias of $V_{dd} = 5V$ was applied to the ring oscillators and the oscillation frequency was measured by AC coupling the V_{dd} bias lead to the input of a spectrum analyzer. The spectrum analyzer displayed the fundamental oscillation frequency as a peak which was usually >6dB above the spectrum analyzer noise floor. Using this method to measure the oscillation frequency only required two leads to be brought off chip for each oscillator: V_{dd} and ground. This method also prevented the need to place an additional load on the ring oscillator for optical or electrical readout.

Each inverter within the ring oscillator was constructed using a 3m width N-FET, and a 6 μ m P-FET, which when simulated, produced approximately equal rise and fall times. Several oscillators were placed on each chip; with, and without 15 μ m x 15 μ m pads and 9 μ m x 9 μ m pads. A value of 14.4fF was calculated for the stray interconnect capacitance between each inverter without any solder bump pads attached. A value of 12.4fF was calculated for the 15 μ m x 15 μ m first level metal pad. The actual circuit layout and wafer parameters for overlap and fringing capacitance parameters were used to make these initial calculations.

Unfortunately, only the 15 μ m pad circuits could be tested with diodes, due to low yield in the attachment of the 9 μ m devices. Four circuits were tested before flip-chip bonding: two circuits with pads, and two circuits without pads. The oscillation frequency was measured at 130.2MHz and 130.7MHz for the two circuits without pads, and 108.4MHz and 108.9MHz for the two with pads. After attaching MQW diodes to the two circuits with pads, the oscillation frequency dropped to 76.7MHz and 78.1MHz. For this last measurement, the MQW diodes were reverse biased using an 8V supply. From this data, the mean gate delays were calculated as 202ps, 244ps, and 340ps for no pads, with pads, and with pads+diodes respectively. The following equations were written:

$$202ps = t_0 + k(14.4fF)/3 \quad (4.9)$$

$$242ps = t_0 + k(14.4fF+12.4fF)/3 \quad (4.10)$$

$$340ps = t_0 + k(14.4fF+12.4fF + C_{diode} + C_{bump} + C_{chip})/3 \quad (4.11)$$

Simultaneously solving the first two equations results in $t_0 = 155ps$ and $k=9.68$. Using these values in the third equation results in $C_{diode} + C_{bump} + C_{chip} \approx 30fF$. This value is significantly smaller than the expected value of $\sim 56fF$, and required further investigation.

First, the sensitivity of oscillation frequency to bias voltage was measured to determine if this could cause the discrepancy. As one would expect, the photodiode bias was found to have little influence on frequency, as long as the diodes remained in reverse bias. The V_{dd} bias was found to change the oscillation frequency at a rate of $\sim 0.02\%/mV$. This translates to a sensitivity in calculated capacitance of $\sim 50aF/mV$. Since the same digital volt meter with better than 10mV accuracy was used for all tests, and special attention was paid to ensure good probe contact, the maximum error in load capacitance calculation due to a metering error would probably be <1fF.

Next, a check was made to ensure all photodiodes were actually making electrical contact. The photodiode bias was set to approximately $V_{dd}-2V$, which caused the photodiodes to become forward biased. The diodes were viewed under magnification with a camera and all were found to luminesce, hence, all were making contact.

The dimensions of the diodes were then measured using a microscope equipped with a measuring stage accurate to better than $1\mu m$. The outer mesa of each diode was found to be only $16\mu m \times 43.5\mu m$, instead of the drawn $22\mu m \times 52\mu m$ indicating that they were severely over-etched. The metal pads could actually be seen protruding from the sides, which meant part of the N-mesa was also etched away.

The diode responsivity to input spot position was then measured using a custom characterization test set designed for S-SEED testing [81]. The results of the test are shown in Figure 13(a), superimposed over an as drawn device topology. Figure 13(b) is a photomicrograph of same photodiode under forward bias. Note the severe over-etching of the device outer mesa when compared to the as-drawn device. Since the fabrication of these devices, a new highly selective etch stop layer was developed [82] which will prevent the over-etching of future device arrays. Also note in Figure 13 the small active area, approximately $8\mu m \times 14\mu m$. Upon talking with the device fabricators, it was discovered that an experimental isolation implant step was added to the process, in order to reduce the diode active area. These devices were inadvertently attached to the ring oscillators. The isolation mask was drawn to produce an active strip $11\mu m$ wide, shown as a dashed line in Figure 13(a). Due to straggle of the implant, the actual active area was reduced to only $\sim 8.5\mu m$ wide.

An FEA simulation was re-run using the measured dimensions, and was found to predict the measured capacitance of $30fF$ when the diode active area was set to $8.7\mu m \times 14.5\mu m$. This compares favorably to the measured active area, thus substantiating the validity of the simulations. The 46% reduction in diode capacitance due to the smaller implant and over-etching may seem attractive, however, both the responsivity ($\sim 0.2A/W$), and active area was far too small for real system applications. Using a corrected implant mask to define an actual active region of $12\mu m$ wide could be beneficial, however, a study done elsewhere found that a mesa etched structure with the same active area as a planar implanted structure, displayed less capacitance, and was attributed to lateral spreading of the depletion layer [83].

Note that a significant amount of active area ($4\mu m \times 17\mu m$) was required in order to make the ohmic contact, which was fabricated to the side of the device optical window. It would be more beneficial to fabricate the ohmic contact within the active region, as is normally done for back illuminated photodiodes. The reason the ohmic contact was placed off to the side, was because the ohmic metal does not have satisfactory reflectance which is an important issue when

the diode is used as a modulator.




Figure 13. (a) Solder bump-bonded MQW Diode response to input spot position contours superimposed over an "as drawn" device topology. An approximate $4\mu\text{m}$ diameter optical spot was used to interrogate the device. The contours represent the normalized response. (b) To scale photomicrograph of same photodiode under forward bias (bright horizontal lines are a result of the image capture). Note the severe over-etching of device outer mesa in comparison to the drawn device. Also note the small active area ($\sim 8\mu\text{m} \times 14\mu\text{m}$) due to an excessive straggle of an isolation implant attempt.

4.2.3 Analysis of Thermal Resistance

When a MQW diode photodetector changes temperature, the absorption band edge and exciton shift approximately $+0.3\text{nm}/^\circ\text{C}$. Thus, if the wavelength of the input light is near this band edge, the amount of photocurrent generated and sent to the receiver will be a strong function of temperature. Thus, it is best to keep the input wavelength sufficiently below the band edge, so changes in temperature have little effect on absorption. However, when the MQW diode is used as an electro-optic modulator, the read-out laser wavelength is intentionally placed near the absorption band edge in order to allow the beam to be modulated using the quantum-confined stark effect (QCSE)[84]. The QCSE causes the absorption peak to shift to a lower energy as the electric field increases. This effect allows the modulator absorption to be a function of applied bias voltage, thus performing the electrical to optical conversion.

To ensure optimum modulator performance, either the case temperature of the smart pixel chip package can be held constant using conditioned air or a thermal-electric cooler, or the MQW modulator bias voltage may be adjusted to compensate for a global temperature change across the chip. Figure 14 shows the modulator response for three possible modes of biasing and a 5V modulation voltage. The effect of a temperature change of 10°C is also shown. In Figure 14(a), a constant bias is set for optimum operation at 25°C . As the temperature increases, the output change in reflectance decreases to nearly zero at 35°C . Figure 14(b) has a constant bias set such that it is optimum for the 25°C to 35°C range. In this case the modulator output reflectance is less than optimum, however, it remains nearly constant, which may be desirable in some system applications. In Figure 14(c), the modulator bias has been adjusted to produce an optimum change in output reflectance for both temperatures. Here the diode output reflectance has been improved over the constant bias mode by 2.1dB and 0.13dB for 25°C and 35°C respectively.

Global temperature-tracking-modulator-bias circuits can run into a potential problem if a thermal gradient, caused by localized heating, occurs across the smart pixel chip. A specialized mount was designed and fabricated to minimize the gradient caused by static localized heating sources on an optoelectronic switch chip.[85] However, if the localized heating is dynamic, then a dynamic temperature gradient could occur.

Dynamic heating can occur within a MQW diode modulator, since the amount of heat generated in the diode is dependent on its state of absorption, and the impinging optical power. The impinging light that is not reflected is absorbed as a photocurrent, which in turn generates heat. Thus, a data dependent heating effect can occur, especially with DC coupled systems which can transmit long strings of '1's or '0's. While suppressed photocurrent MQW modulators have been demonstrated to reduce the generation of heat, [86] the current solder bump process has been refined to only attach one type of device. Multiple attachment of non-interspersed diodes has only recently been demonstrated [87]. For this reason, the thermal conduction of the solder bumped MQW diode modulator islands to the silicon substrate is an important parameter.

Section 4.3 investigates the thermal conduction of the solder bumped MQW diode modulator islands. Section 4.3.1 first estimates the thermal resistance using empirical formulae. Section 4.3.2 presents the results of a three-dimensional finite element analysis simulation. Finally, Section 4.3.3 presents measured results of the thermal effect in solder bumped diodes to silicon circuits.

Estimated Thermal Conduction

The conduction of heat from an active device to an ultimate heat sink may be characterized by a thermal resistance parameter, which is analogous to electrical resistance. The thermal resistance in $^{\circ}\text{C}/\text{W}$ is defined [88] by

$$R_{\text{TH}} = \{DT\}/Q, \quad (4.12)$$

where DT is the temperature drop in $^{\circ}\text{C}$, and Q is the heat flow in Watts. If the thermal resistance is known, the total temperature rise of a device may be calculated by multiplying the thermal resistance by the electrical power dissipated in the device. The total thermal resistance may be divided up into several contributing components that act in series or parallel to each another. The resulting series/parallel resistor network may then be solved to determine the total thermal resistance.

The thermal resistance of the MQW modulator diode island to the silicon substrate was estimated using the simplified model shown in Figure 15. Within this model, the silicon substrate was assumed to be the heat sink, and the analysis focused on the temperature rise of the diode active area, relative to the temperature of the silicon substrate. Although the diode island would actually consist of GaAs/AlGaAs layer structure shown in Figure 15, it has been simplified to be a block of GaAs. This assumption leads to a lower predicted R_{chip} thermal resistance, since the thermal conductivity of AlGaAs is less than that of GaAs[89]. Figure 15 depicts two thermal paths and six thermal resistance components within the thermal network used to model heat conduction. Also shown in Figure 15 are the formulae used to estimate each component. The R_{SiO_2} component represents the thermal resistance of the silicon dioxide and dielectric layers that exists between the bonding pad and silicon substrate. Note that additional layers of dielectric may also exist under the bonding pad, depending on what level metal is used. To simplify the calculations, a stand-alone pad fabricated in first level metal was assumed, and thermal spreading from pad to substrate was ignored. In actual circuits, the connecting metals would serve to further reduce the thermal resistance, thus, the result here should be considered a conservative value. The R_{bump} components represents the solder and interconnection metals. R_{bump} may be further divided into the thermal resistance contribution on the actual metals diagramed in Figure 15. The PbSn solder used in the sample presented here had the lowest thermal conductivity of the metals comprising the bump and thus dominates the bump thermal resistance. To simplify the calculations, the effect of the other metals will be ignored. The R_{chip} component in Figure 15 represents the thermal resistance of the portion of the GaAs chip that bridges the two pads, and was modeled as a slab of GaAs between the two bump midpoints. The R_{spread} component represents the spreading resistance from the diode active area to the solder bump. The spreading resistance is a result of assuming the input light is a small diffraction limited spot, which generates heat in an area smaller than the area of the solder bump. This situation is shown in Figure 16.

With the exception of R_{spread} , the formulae used to estimate the components shown in Figure 15 assumed columnar heat flow. This assumption allowed the conductive thermal resistance to be described using:

$$R_{\text{TH}} = l/kA, \quad (4.13)$$

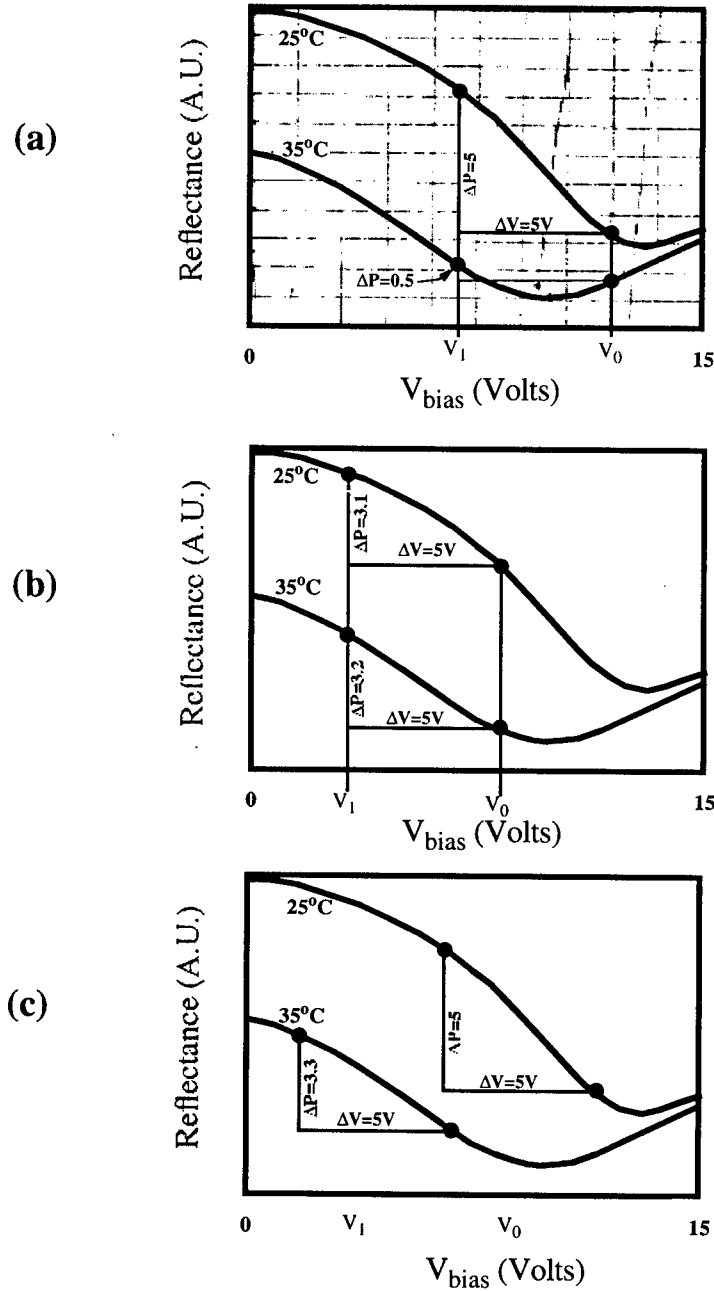


Figure 14. DC plots of MQW diode reflectance vs. applied bias data for two case temperatures of 25°C and 35°C. Three possible modes of operation are depicted: (a) Constant V_{bias} optimized for 25°C, (b) Constant V_{bias} optimized for the 25°C-35°C range, and (c) Temperature tracking V_{bias} . A ΔV of 5V was used in all cases, and diode output reflectance (ΔP) is shown as arbitrary units. Temperature tracking in (c) improves diode output reflectance by 2.1dB and 0.13dB for 25°C and 35°C respectively, over the constant bias mode shown in (b).

where A is the cross-sectional area in m^2 , l is the length of the thermal path in meters, and k is the

thermal conductivity of the material in W/(mK). For most materials, the thermal conductivity decreases only slightly over the temperature range of interest (300°C-360°C), so k is assumed to be a constant. The spreading resistance for R_{spread} may be estimated using the empirical relation developed by Kennedy:

$$R_{\text{spread}} = H_2 / (kpr), \quad (4.14)$$

where r is the radius of the heat source, and H_2 is the empirical spreading resistance factor. The parameter H_2 is a function of the diameter and location of the heat source and sink, and is equal to 0.30, 0.31, and 0.32 for pad dimensions of 5 μm , 10 μm , and 15 μm . These values assumed a 4.3 μm diameter uniform heat source located 0.74 μm from the bump metal, which approximates a typical spot produced by the Bell Labs photonic switch optics centered within the device active area. Table 8 summarizes the results of using the formulae in Figure 15 to calculate the thermal resistance for islands with pad dimensions of 5 μm , 10 μm , and 15 μm . The total resistance (R_{total}) was found by solving the series/parallel combination of:

$$R_{\text{total}} = (R_{\text{spread}} + R_{\text{bump}} + R_{\text{SiO}_2}) \parallel (R_{\text{chip}} + R_{\text{bump}} + R_{\text{SiO}_2}). \quad (4.15)$$

THERMAL THERMAL RESISTANCE (°C/W)

RESISTANCE COMPONENT	d=15 μm h=10 μm	d=15 μm h=5 μm	d=10 μm h=10 μm	d=10 μm h=5 μm	d=5 μm h=5 μm
R_{SiO_2}	4,400	4,400	10,000	10,000	40,000
R_{bump}	1,230	610	2,780	1,390	5,560
R_{chip}	9,300	9,300	7,900	7,900	5,480
R_{spread}	850	850	880	880	910
R_{total}	4,500	4,160	8,200	7,500	24,300

Table 8. Estimated thermal resistance for several geometries.

It is evident from table 8 that the dominant contributor to the thermal resistance is the R_{SiO_2} component, due to its low thermal conductivity of only $\sim 1\text{W/mK}$. This is especially the case for small pad dimensions which have a small cross-sectional area. It is anticipated that in an actual system, the R_{SiO_2} component would be reduced by other interconnecting metals, however, the trend of increasing thermal resistance for small pads would still be a factor to consider.

As a comparison, the maximum thermal resistance may be estimated for the monolithic FET-SEED technology as the power dissipated in a circular disk above an infinite substrate using the relation [91]

$$R_{\text{TH_monolithic}} = 1 / 4 k r, \quad (4.16)$$

where $k=52\text{W}/(\text{mK})$ is the thermal conductivity of GaAs, and r is the radius or the optical input spot. Using the same optical spot size as the hybrid example ($r=2.15\mu\text{m}$), the maximum thermal resistance would be $2,200^\circ\text{C}/\text{W}$ in the monolithic technology. Thus, the thermal resistance in the monolithic technology is primarily a function of the spot size, although thermal resistance could be further reduced by thinning the substrate.

If the GaAs substrate were retained in the hybrid process, it would serve to act as a heat spreader. Additional solder bumps, or epoxy, could be used with the sole purpose to sink the heat. Potential problems with retaining the substrate include the coefficient of thermal expansion mismatch between the large GaAs and Si chips, and the requirement of additional substrate vias to allow the 850nm light to pass through to the MQW diodes, since the GaAs substrate is absorptive at the operating wavelength. Further investigation of effect of the substrate on thermal resistance is required.


The source labeled P shown in Figure 15 represents the power dissipated in the diode active area. The amount of power dissipated depends on the modulator output state (high or low), and the amount of impinging optical power. Assume the two states of the output driver place a 2V and 7V reverse bias across the modulator operating at I_1 . In the I_1 mode of operation, the input wavelength is set such that the exciton absorption peak occurs when the modulator is under high reverse bias. Thus, the 2V bias represents an output logic high, with a diode responsivity of $\sim 0.2\text{A}/\text{W}$. The 7V bias represents an output logic low, and a diode responsivity of $\sim 0.6\text{A}/\text{W}$. If the read beam optical power is $500\mu\text{W}$, the difference in power dissipated between each state would be $(500\mu\text{W})[(7\text{V}+V_{bi})(0.6\text{A}/\text{W}) - (2\text{V}+V_{bi})(0.2\text{A}/\text{W})]$. Assuming a built-in diode potential of $V_{bi}=1\text{V}$, the power difference would be approximately 2mW . Note, the majority of this power would be dissipated within the diode active area, with only a small percentage ($\sim 2\%$) dissipated in the resistive P-layer.

$d=15\mu\text{m}, h=10\mu\text{m}$	$d=15\mu\text{m}, h=5\mu\text{m}$	$d=10\mu\text{m}, h=10\mu\text{m}$	$d=10\mu\text{m}, h=5\mu\text{m}$
9.0°C	8.3°C	16°C	15°C

Table 9. Estimated temperature differential between two output logic states for several geometries.

Table 9 summarizes the calculated temperature differential for the two logic states of the modulator for two different pad dimensions and bump heights. Note that for the MQW structure used in this modulator, the exciton was measured to shift at a rate of approximately $-0.3\text{V}/^\circ\text{C}$, therefore, a 9°C rise in temperature would shift the exciton by approximately -2.7V . This shift could significantly decrease the differential output power level depending on the initial bias condition (see Figure 15). This temperature differential would be most problematic for systems which allow long strings of '1's or '0's, since the thermal time constant to change temperature due to the absorption state would be relatively long compared to bit intervals at 100Mbps or greater data rates [92]. If the data was encoded with a 50% duty cycle, then an average steady-state temperature would be reached, approximately equal to the mean of the two states. Note that even when encoded data is used, the temperature of the diode would still be dependent on the amount of impinging optical read-out power, hence, uniform optical read beams would be required.

The thermal analysis in this chapter assumed the underlying silicon substrate acts as a constant



$$R_{SiO_2} = \frac{t_{SiO_2}}{(k_{SiO_2})d^2} + \frac{t_{di1}}{(k_{di1})d^2} + \frac{t_{di2}}{(k_{di2})d^2} + \dots$$

$$R_{bump} = \frac{h}{(k_{Solder})d^2}$$

$$R_{chip} = \frac{2d}{(k_{GaAs})(t_{chip})(d + 8\mu m)}$$

$$R_{spread} = \frac{H_2}{(k_{GaAs})\pi r}$$

Assumed Thermal Conductivities

$$k_{SiO_2} = 1\text{W/mK}, \quad k_{din} = 1\text{W/mK}, \quad k_{solder} = 36\text{W/mK}, \quad k_{GaAs} = 52\text{W/mK}$$

Figure 15.

temperature heat sink. While the thermal conduction of silicon is very good ($k=163\text{W}/(\text{m K})$), the active circuitry (i.e. the receiver, logic, and transmitter) would also dissipate static and dynamic power and should be considered when designing a smart pixel system. The magnitude and effect




Figure 4.7. Diagram of solder bump metal used in the hybrid process




Figure 16. Diagram showing the thermal spreading from the diode active area to the solder bump metal.

of these additional heat sources would be highly dependent on the circuit implementation, and will not be analyzed here.

Simulated Thermal Conduction

The model used to estimate the thermal resistance in the previous section assumed there was no epoxy surrounding the GaAs island. The fabrication process can include a step to remove the epoxy, however, the epoxy can also be left on the chip, providing that it does not interfere with the attachment of the electrical bond wires on the chip perimeter. In order to determine the effect of epoxy on thermal resistance, a finite element analysis tool was required, due to the complex geometry and boundary conditions of the epoxy surrounded islands.

A finite element analysis was first run on a 15 μ m pad, 10 μ m high bump device without epoxy, in order to check the validity of the approximations. This resulted in a R_{total} of 4,670°C/W, which compared favorably with the approximation of 4,500°C/W. When the epoxy was added to this structure, R_{total} was reduced to 3,310°C/W, hence, even with the low thermal conductivity of only 1W/(m K), the epoxy has dramatically reduced the total thermal resistance by ~30%. A simulation of a 5 μ m bump height with 15 μ m pads resulted in a R_{total} of 3,000°C/W, which was ~28% less than the estimated value for the epoxyless structure.

4.3.3 Measured Thermal Conduction

An attempt was made to determine the thermal resistance by measuring the exciton shift as a function of increasing modulator optical input power. The increase in optical input power would result in a rise of photocurrent, and hence, a temperature rise of the device active area. The temperature rise would cause the location of the exciton to appear at a lower applied bias voltage. Two chips were tested. The first was a photonic switch chip, which contained MQW diodes connected to Si circuitry. The second chip, which became available at a later date, contained solder-bumped MQW diodes with independent contacts brought out to wire-bonding pads.

The photonic switch test chip had receivers with two series connected input photodiodes. One photodiode was connected with its cathode to an amplifier input and its anode to an external bias supply lead. The other diode was connected with its anode to the same amplifier input, and its cathode to a different external bias supply lead. To test a single diode response using this circuit, two optical beams were required. One high power beam (~2mW) was used to effectively short the diode connected to ground. The other beam served as the interrogation beam for the other diode.

A test was performed by sweeping the photodiode bias voltage around the position of the exciton for a fixed input wavelength while monitoring the device photocurrent and reflectance. The optical input power was incrementally increased, and the location of the exciton in applied bias voltage was noted. The data from this test is shown plotted in Figure 17. The line shown in Figure 17 is an estimation of the response. Note that if only the thermal effect was present, the exciton would appear to shift to a lower applied bias voltage for increasing optical powers. The inflection point within the data may be attributed to the influence of two other factors. One factor is the IR voltage drop across the parasitic series resistance, and the other factor is the MQW diode's forward bias characteristic. Both of these would contribute to a perceived positive shift in the location of the exciton with respect to the applied bias voltage.

The parasitic series resistance contributes a perceived linear positive shift (i.e. higher bias voltage) in exciton location, since as the input power is increased, the photocurrent increases and causes a voltage drop across the resistive elements. Thus, a larger external bias is required to develop the same potential across the diode by an amount equal to $I_{\text{ph}} R$, where I_{ph} is the resulting photocurrent, and R is the total series parasitic resistance. The dominant contributors to the parasitic series resistance were expected to be the bulk resistivity of the P-layer (~410 Ω) within each diode and the N-ohmic contacts (~30 Ω each) and the metal traces (~40 Ω total). These are shown as resistors in the circuit drawn within the plot of Figure 17. The total resistance for the series connected pair was estimated to be ~1k Ω by using the typical bulk resistivities and the as-drawn device geometry for d=15 μ m pads.




Figure 17. Measured location of exciton in applied bias voltage as a function of optical input power for a series connected diode pair. The line shown is an estimation of the continuous characteristic. The circuit used to take the data is also shown. The inflection point labeled Q, is attributed to the voltage drop across the other forward biased MQW diode, which has canceled the shift due to heating.

The forward bias characteristic of the other MQW diode also contributes a perceived non-linear positive shift of the exciton. The voltage drop across a forward bias diode (V_{diode}) is also a function of current, and may be described as [93]

$$V_{diode} = \frac{hkT}{q} \ln \left(\frac{I}{I_s} \right) \quad (4.17)$$

where I is the current through the diode, I_s is the reverse saturation current of the diode, T is the

absolute temperature, h is the diode ideality factor, q denotes electronic charge, and k is the Boltzmann constant. Using the above equation, the change in diode forward bias voltage drop due to a change in current (DI) from a reference current (I_{ref}) would be equal to:

$$DV_{diode} = \frac{hkT}{q} \ln(1 + DI/I_{ref}) \approx \frac{hkT}{q} DI \quad (4.18)$$

An attempt was made to extract the MQW diode forward biased characteristic and parasitic series resistance. The MQW diode characteristic was measured by sweeping a positive voltage on the $-V_{bias}$ lead shown in the circuit in Figure 17, while the amplifier V_{dd} bias lead was held at ground potential. The conduction path included a parasitic forward biased diode within the feedback amplifier on the CMOS chip. The data from this measurement included 43 other similar circuits in parallel, since this was a parallel array. After normalizing for the 43 circuits, and subtracting the effect of the CMOS diode, the following parameters were extracted for a single MQW diode: $I_s = 10^{-14}$ A, $R_s = 2.8$ kW. The value for I_s is typical, however total parasitic series resistance for two diodes of $2R_s = 5.6$ kW is significantly larger than the previously estimated 1 kW. Later testing of an unbonded GaAs MQW diode array from the same wafer revealed a high series resistance. The high resistance was attributed to an improper doping of the N-mesa which caused a high ohmic contact resistance.

The location of the inflection point in Figure 17 indicates that the negative shift of the exciton bias point caused by the thermal resistance has been offset by the combined positive shifts of the series resistance and forward biased diode. This situation may be written as:

$$V_{inflection} DIgR_{TH} = DIR + \left(\frac{hkT}{q}\right) \ln(1 + DI/I_{ref}) \quad (4.19)$$

where DI is the difference in photocurrent between a reference current (I_{ref}) and the current at the inflection point, $g = -0.3$ V/°C, and R is the parasitic series resistance. This equation may be used to solve for the thermal resistance R_{TH} . In this case, the reference current (I_{ref}) was taken as the first data point of ~ 19 μ A. The location of this inflection point was estimated from the data shown in Figure 17 to be at $V_{inflection} = 4.89$ V and $I_2 = 140$ μ W (67 μ A). Using this data in the above equation resulted in a $R_{th} = 5,000$ °C/W, which is greater than the simulated value of $\sim 3,000$ °C/W for a structure surrounded by epoxy. The increase was attributed to the fact that the simulation model assumed first layer metal was used for the diode bonding pads. The diode bond pads on the test chip had been fabricated in second layer metal, which added an additional dielectric layer under the bond pad. The specified thickness of the dielectric layer was 0.8 μ m. Assuming a thermal conductivity of 1 W/m k, this would increase R_{SiO_2} an additional 3,500 °C/W. Including this additional resistance in the thermal network, the total R_{th} was re-calculated to be 6,500 °C/W for an epoxy-less structure. Assuming the same 30% reduction for the structure with epoxy, a $R_{th} = 4,500$ °C/W would be expected. This value agrees to one significant digit with the measured value of 5,000 °C/W.

A second test chip, which became available at a later date, contained independent MQW diodes and was also tested. This chip was mounted in a custom package with a thermo-electric-cooler and thermister. The following parameters were extracted from the forward bias characteristic curves: $I_s \approx 10^{-13}$ A, $R_s = 3.3$ kW. The high parasitic series resistance of 3.3 kW was once again

attributed to an improper doping on the N-mesa leading to the high ohmic contact resistance, even though the MQW diodes came from a different wafer. The location of the exciton absorption peak in applied bias volts was measured, and is shown in Table 10 for several wavelengths and two temperatures. The peak was located by plotting the derivative of the measured diode reflectance and noting the zero crossing using an hp4145B semiconductor parameter analyzer.

WAVELENGTH EXCITON LOCATION IN

(nm)	APPLIED BIAS VOLTS	
	@25°C @35°C	
855.5	15.1	12.0
854.3	13.75	10.8
853.3	12.9	9.9
852.4	12.3	9.5

Table 10. Measured location of exciton absorption peak as a function of wavelength for two case temperatures of 25°C and 35°C.

The change of the exciton location in applied bias voltage for a change in temperature may be modeled as a first order linear approximation of $-0.3\text{V}/^\circ\text{C}$. This sensitivity parameter will be referred to as γ .

The shift of the exciton absorption peak as a function of optical input power was measured for two wavelengths (843.6nm and 852.2nm) and is shown plotted in Figure 18. Note the absence of the inflection point, as seen in Figure 17, since the circuit no longer contains the second series connected MQW diode. The sensitivity for both wavelengths is approximately $-7\text{mV}/\mu\text{W}$ over input powers of 25 to $300\mu\text{W}$ (10.8dB). Note that this sensitivity includes the effect of the parasitic series resistance (R_s). Using the extracted value of $R_s=3.3\text{k}\Omega$ and a mean peak responsivity of 0.4A/W , the positive shift due to R_s would be $(0.4\mu\text{A}/\mu\text{W})(3.3\text{k}) = +1.2\text{mV}/\mu\text{W}$. Thus, the sensitivity attributed to thermal heating alone would be approximately $-8.2\text{mV}/\mu\text{W}$. This sensitivity parameter will be referred to as α .

The thermal resistance may be calculated using the sensitivity parameters γ , α , and the measured mean exciton location V_{exciton} , and the peak responsivity S_{exciton} to be:

$$R_{\text{th}} = \alpha / (\gamma V_{\text{exciton}} S_{\text{exciton}}) \quad (4.20)$$

Using the values of $a=-8.2\text{mV}/\mu\text{W}$, $\gamma=-0.3\text{V}/^\circ\text{C}$, $V_{\text{exciton}}=8.3\text{V}$, and $S_{\text{exciton}}=0.4\text{A/W}$: $R_{\text{th}} = 8,200^\circ\text{C/W}$. Recall that the simulated thermal resistance value for this epoxied structure from device active area to Si substrate was $\sim 3,000^\circ\text{C/W}$. The Si substrate, die bonding material, and copper heat extractor would add to the overall resistance, however, it would be expected to be on the order of 100°C/W or less. Further investigation found that the diodes were bonded to pads made of third layer metal. The analysis in Section 4.3.1 and 4.3.2 assumed that the first level metal was used for pads. When using the third layer metal, two additional layers of dielectric are

present between the pad and the SiO_2 . The thickness and thermal conductivity of the two additional dielectric layers for this chip were unknown, however, the value of $8,200^\circ\text{C/W}$ seems reasonable in the previous progression from 3000°C/W to 5000°C/W for first and second layer metal pads respectively.




Figure 18. Measured location of exciton in applied bias voltage as a function of optical input power with wavelength as a parameter for a stand-alone solder-bumped MQW diode. The solid line and dashed line represent 843.6nm and 852.2nm respectively. The circuit used to take the data is also shown. The resistive elements represent the parasitic resistance of the interconnection path.

It is interesting to note that for a stand-alone diode, an optimum value of parasitic series resistance (R_s) could be chosen to completely offset the temperature effect. The optimum resistance would be determined by setting

$$R = \gamma R_{TH} V_{exciton} \quad (4.21)$$

For example, using the previously measured parameters of $\gamma = 0.3 \text{ V}/^\circ\text{C}$, $R_{TH} = 8,200^\circ\text{C}/\text{W}$, and $V_{exciton} = 8.3 \text{ V}$, the optimum series resistance would be $R_s = 20 \text{ kW}$. This phenomenon is shown experimentally in Figure 19. Figure 19(a) shows the MQW diode reflectance vs. applied reverse bias voltage for two input powers of $25 \mu\text{W}$ and $125 \mu\text{W}$. The derivative of the response is also shown as a dashed line, with the zero crossing indicating the location of the exciton absorption peak. A difference in the location of the exciton of 1.1 V between the two optical input powers is clearly shown. Since the parasitic series resistance of the diode is 3.3 kW , an additional $20 \text{ kW} - 3.3 \text{ kW} = 16.7 \text{ kW}$ resistor was added in series, with the results shown in Figure 19(b). It may be seen that the exciton absorption peak occurs at approximately the same bias voltage for both optical input powers, thus the series resistance has offset the effect of heating. While this is a novel way to offset the temperature effect, its applicability to high data rate modulators is limited. This is because the high series resistance combines with the parasitic capacitance to slow down the achievable modulation rate.

Temperature Compensated Modulator Bias Circuit

Global changes in optoelectronic chip temperature can induce changes in the optical output characteristics of either vertical cavity surface emitting lasers (VCSELs) [95] or MQW modulators [96]. Conditioned air or thermo-electric-coolers are typically used to control the temperature of the optoelectronic chip packages. These techniques tend to have a slow response time and usually contain some thermal oscillation, as the servo loop between the sensor and heating/cooling element adapts to temperature changes.

By placing a sensor on the surface of the chip, and controlling a parameter such as bias voltage or current, a faster response may be obtained. Having a sensor on the chip surface can also compensate for dynamic changes in chip temperature due to changes in on-chip power dissipation. An external sensor is unable to compensate for temperature changes of this sort, because of the thermal resistance and thermal time constant that exists between the chip surface and package case.

Temperature-independent biasing is a common technique used in analog integrated circuit design [97]. For example, temperature-independent current sources have been demonstrated using Zener or band-gap referencing [97]. It is also possible to construct a proportional-to-absolute-temperature (PTAT) current reference. For example, a CMOS PTAT current reference was used within an LED driver to provide a temperature compensated bias current [98]. A temperature compensated bias circuit for a MQW diode modulator requires a very strong negative temperature coefficient. The required temperature coefficient (γ) for the devices reported in Section 4.3 was approximately $-300 \text{ mV}/^\circ\text{C}$. As seen in Figure 14, the optimum bias point would change approximately at γ , or a slightly higher rate. For an applied bias of 6 V , $300 \text{ mV}/^\circ\text{C}$ represents a change of $50,000 \text{ ppm}/^\circ\text{C}$, which greatly exceeds the $2400 \text{ ppm}/^\circ\text{C}$ capability of the circuit in Reference [98]. Another disadvantage of the circuit described in Reference [98], is that it required an on-chip resistor with a low temperature coefficient to produce the required PTAT.

This section describes a novel technique where two programmable PTAT current references are combined to produce a current reference with the required large temperature coefficient to either

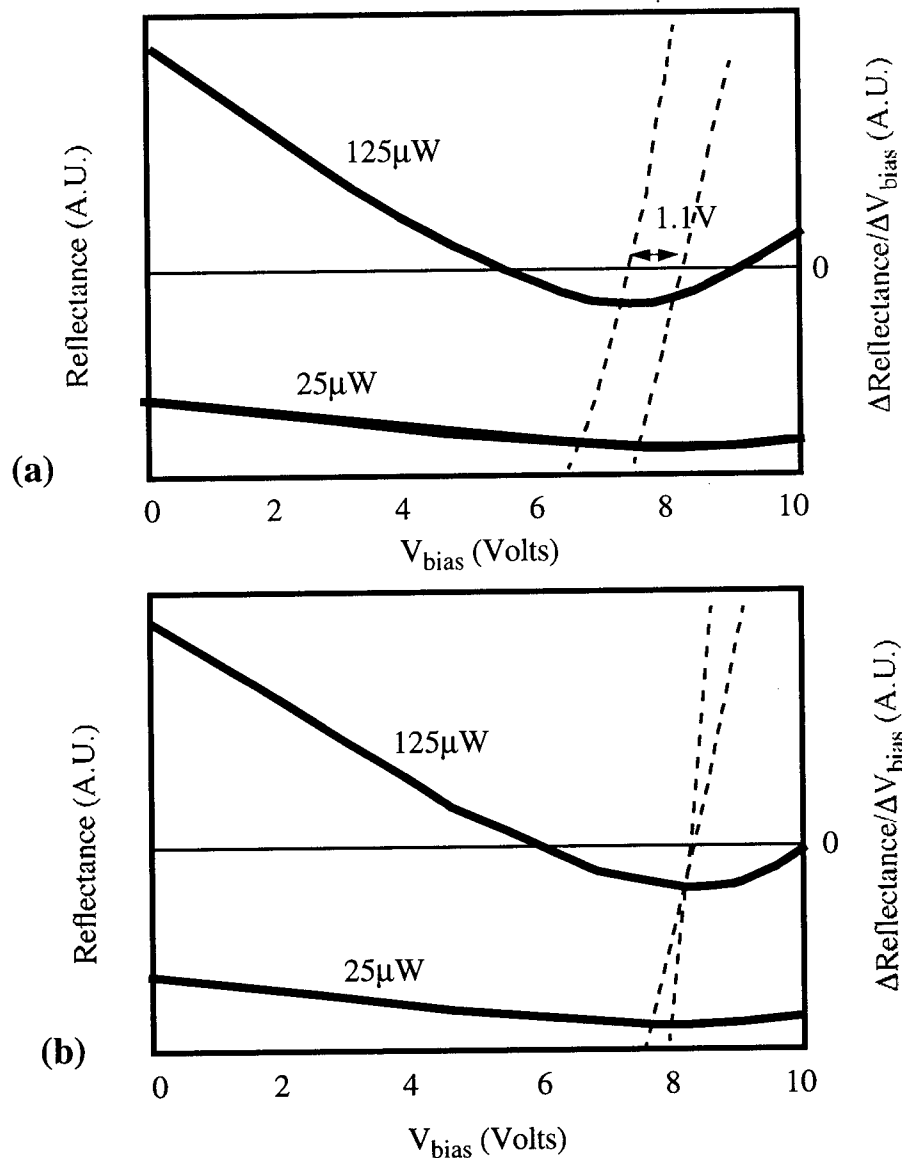


Figure 19. (a) MQW Diode reflectance vs. applied reverse bias voltage for two optical input powers of 25 μW and 125 μW . The derivative is also shown plotted, and depicts a 1.1V difference between the location of the exciton for the two input powers. (b) Same scale plot as in (a), except an additional 16.7k Ω resistor was added in series to counteract the offset due to heating. Here the location of the exciton in applied bias voltage is nearly independent of applied optical input power.

bias VCSELs or MQW diode modulators.

A block diagram of the PTAT current referenced modulator bias supply circuit is shown in Figure 20. All blocks within the dotted line are contained on the CMOS chip. The on-chip circuit outputs a current with a strong negative temperature coefficient that is transformed to a voltage by an external resistor. The off-chip resistor requires a low temperature coefficient, or must be kept at a constant temperature, to avoid influencing the response. The voltage across the resistor is amplified by an external DC power supply. This amplifier has a fixed gain, and supplies the chip with a

PTAT modulator bias voltage. Although a single supply is shown, a negative PTAT modulator bias voltage could also be produced to supply bias for differential modulator configurations. The reason an external supply is used, instead of an on-chip regulator, is because the modulator bias would typically need to be adjusted to a value greater than 12V, which would exceed the breakdown voltage of the CMOS FETs.

Note that for a VCSEL PTAT bias circuit, the polarity of the inputs to the on-chip current summation point would be reversed in Figure 20. This would provide a strong positive temperature coefficient, usually required by VCSELs.

The novelty of the circuit is the strong negative temperature coefficient that may be obtained by subtracting the output of a constant current source from the output of a current reference with a negative temperature coefficient. Flexibility is added to the circuit by making both current references, as well as the output current mirror programmable. Programmability is required in order to match the particular characteristics of the CMOS wafer to that of the modulators that are later attached to the circuit. A 3-bit fusible-link trim is performed on a programming resistor to set the output of each current reference, although this could have been accomplished using on chip pass-transistors and field-programmable control logic.

Circuits were designed, and simulated using an AT&T proprietary simulation tool similar to SPICE,[99] and device models for the MOSIS 0.8 μ m CMOS process. The simulation allowed the optimization of the programming resistors and active device sizes. The combination of the on-chip circuit and external resistor was to provide an output voltage from 1.5V @ 25°C to 0.7V @ 50°C corresponding to a coefficient of $(0.7-1.5)/25=-32\text{mV}/^\circ\text{C}$. This voltage would then be amplified by the external DC power supply amplifier with a gain of ten to provide a modulator bias of 15V and 7V for 25°C and 50°C respectively.

Test circuits were fabricated in 0.8 μ m CMOS, however, initial testing revealed a problem with the N-wells for the P-MOS transistors used in the current sources. The N-wells of the cascade transistors used in the current sources needed to be separated. Although specific wells and contacts were drawn separately for each transistor, I was unaware that the logical layout editor inadvertently ties the N-wells together. This feature does not present a problem in digital logic circuits, however, in the analog current reference circuits it caused an effective $\sim 2\text{k}\Omega$ N-well resistor to be in parallel with the desired $20\text{k}\Omega$ resistor. Thus the operating currents were off by more than a factor of 10.

Despite this setback, a PTAT voltage bias circuit was partially demonstrated using the output from an individual PTAT current reference test circuit. A Si test chip was mounted in a custom package containing a thermo-electric-cooler and thermister. The output from a PTAT current reference was brought off chip and connected to a $2.2\text{k}\Omega$ resistor tied to the circuit's Vdd. The reason the resistor was connected to Vdd was because the test circuit output acted as a current sink. This produced an output voltage (V_{out}) with a positive temperature coefficient with respect to ground. This voltage was fed to the inverting input of a fixed gain differential amplifier, which had its non-inverting input connected to a constant voltage source. The constant voltage source was used to subtract off the large offset current that was caused by the shorted N-wells. The output of the differential amplifier was fed to a high voltage amplifier with adjustable gain. Table 11 shows data

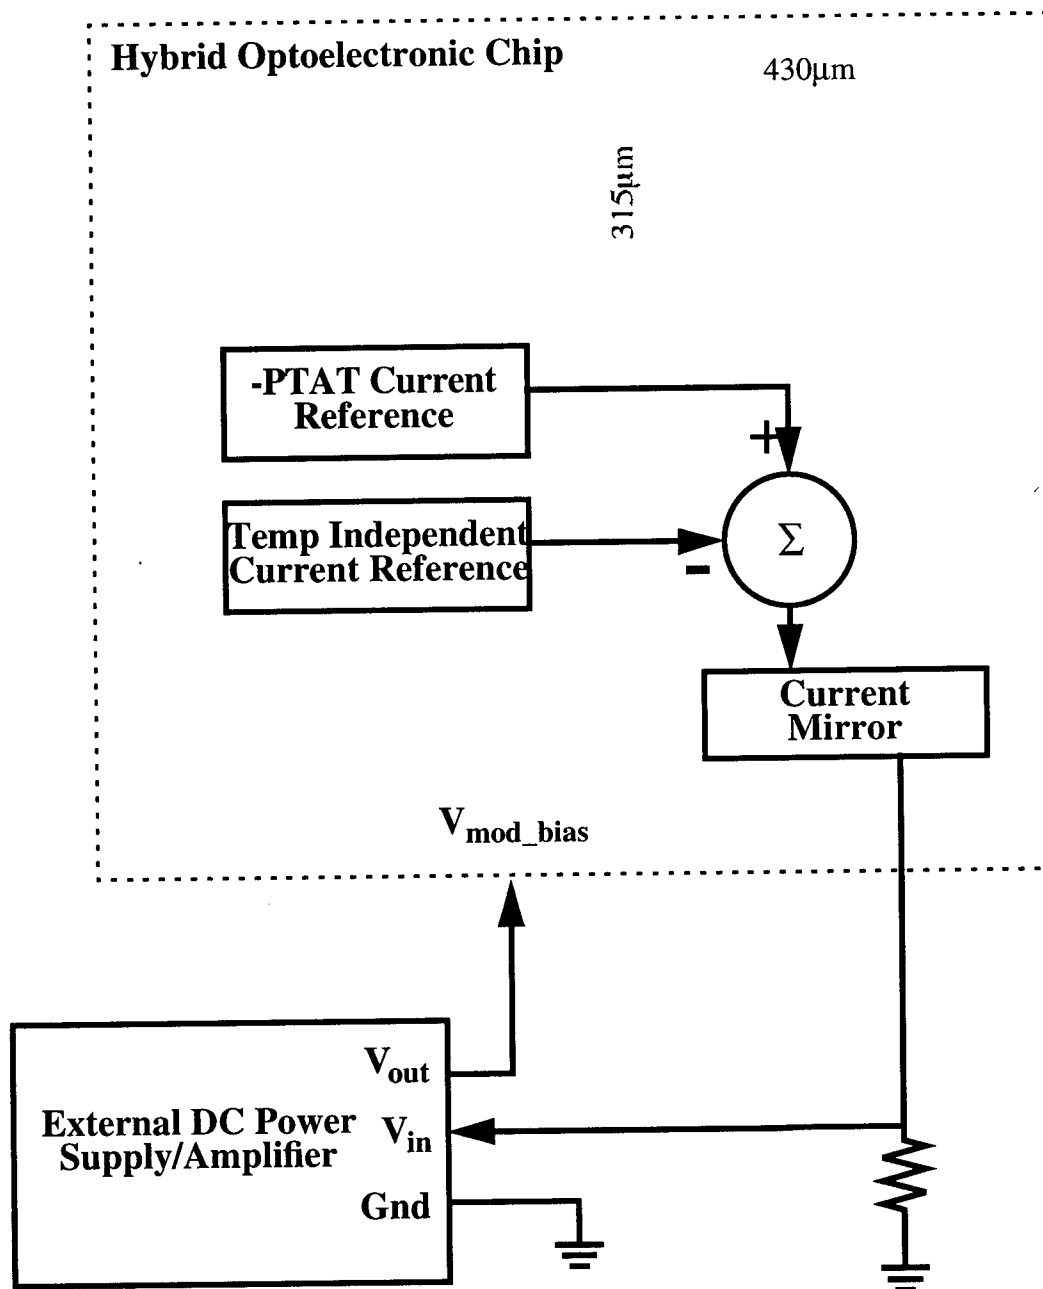


Figure 20. Block diagram of the PTAT current-referenced modulator-bias-supply circuit. All blocks within the dotted line are contained on the chip. A photograph of the $0.8\mu\text{m}$ CMOS chip is also shown.

taken for temperatures of 25°C to 50°C .

Temperature (°C)	I_{out} (μA)	V_{mod} (V)
25	815	14.8
30	800	12.8
35	788	11.2
40	775	9.5
45	765	8.1
50	754	6.7

Table 11. Data from the -PTAT current reference test circuit

As evident in Table 11, the actual output current was much greater than the design output current of approximately $50\mu A$ at $25^{\circ}C$. To produce this large current, the transistors in the current mirrors were operating with a large gate-to-source voltage which also caused a strong dependence of output current on the applied supply voltage. The cause of the problem (shorted wells) has been fixed and submitted to a $0.35\mu m$ wafer run.

4.2.4 Summary and Conclusions

The effect of solder bump geometry used in flip-chip-bonded GaAs SEED photodetectors on Si CMOS was analyzed theoretically and compared to simulated and measured results. Both capacitance and thermal issues were investigated.

When using the MQW diode as an input photodetector, it is desirable to reduce the front-end capacitance for high data rate operation. Formulae were presented to estimate the capacitance of the MQW diode island structure. The total capacitance was found to be dominated by the MQW diode active area and bonding pad. Numerical simulations were also carried out, and indicated that the estimations overestimated the capacitance of the structure by 10%-15%. A conservative estimated value is desirable when designing a circuit, to allow some margin for variations in parameters. The simulations also indicated that the currently used common cathode configuration results in a slightly higher capacitance (~8%) than a common anode configuration. This was found to be especially true for small bump heights, and is attributed to the GaAs chip component adding more significantly to the input capacitance in the common cathode case.

When using the MQW diode as an output modulator, the power dissipated could be on the order of a few mW's, hence the thermal resistance of the MQW diode islands becomes an important parameter. The thermal resistance of the island was found to be dominated by the silicon dioxide and dielectric layers under the silicon CMOS bonding pads. For reasonably sized pads of $15\mu m$, the resulting drop in output contrast would be negligible, however, for smaller pads with high input powers and DC coupled data, the exciton could shift by more than 2V. Since the data could be encoded to remove any DC component, the actual temperature would settle to a steady-state value due to the long thermal time constant. Thus, a pattern dependent thermal shift would not be

expected within photonic switching systems that transport encoded data. The power of the optical read beams, however, would need to be well controlled, as a change in read beam power could induce a thermal shift causing a substantial drop in the change in reflectance.

Although use of higher level metals for bonding pads is desirable to allow the attachment of diodes over active circuitry, the thermal resistance was found to increase for diodes bonded to higher level metals. This is caused by the additional dielectric layers under the pads. Initial test results indicated an increase of the thermal resistance by a factor of 1.6 and 2.6 for second and third level metal respectively, compared to first level metal.

The effect of the epoxy, used within the hybrid fabrication process, was analyzed on both thermal resistance and capacitance. Simulations show that the epoxy can lower the total thermal resistance by ~30% for a 15 μ m pad structure bonded to first layer metal. This was a surprising result, given that the epoxy only had a thermal conductivity of only 1°C/W. The epoxy was found to slightly increase the capacitance of the structure (~5%), due to the relative dielectric constant of 2.9. Since the epoxy may either be left on or removed, a trade-off exists between better thermal conduction and lower capacitance. The following list summarizes some of the desirable changes to the current common cathode hybrid process, or circuit design, to reduce the thermal resistance and parasitic capacitance.

- Use an epoxy with a high thermal conductivity and low dielectric constant.
- Use a large diode cathode bond pad connected to a wide V_{bias} bus.
- Use first layer metal for diode cathode bond pad, a higher level for anode.
- A small diode anode contact that scales with the silicon technology.
- Place the photodiode ohmic contact within the optical window.

The MQW diode parasitic series resistance was found to help offset the exciton shift due to heating caused by an increase in optical read beam power. As the optical input power increases, the voltage drop across the parasitic series resistance also increases, thus requiring a larger applied bias. The same increase in optical power also results in increased power dissipation, which raises the device temperature and lowers the exciton bias point, thus, it is possible for the two effects to cancel one another. For the diodes tested here, a series resistance value to produce satisfactory cancellation for optical input powers of up to 500 μ W would be approximately $\gamma V_{\text{exciton}}$ times the thermal resistance, where γ is the sensitivity of the exciton to a change in temperature in V/°C, and V_{exciton} is the mean location of the exciton in applied bias voltage. This cancellation technique is novel for DC operation, however, when operating using high speed encoded data, a steady-state temperature would be reached, and the series resistance would only serve to slow down the output driver.

The concept for a novel modulator bias circuit that provides a proportional to absolute temperature bias voltage was described. This circuit could be used to cancel the effects of global changes in optoelectronic chip temperature, and reduces the amount of tweaking required to obtain optimum modulator performance. Flexibility was added to the circuit by making the current reference

programmable. Programmability was required in order to match the particular characteristics of the CMOS wafer to that of the modulators that are later attached to the circuit. The novelty of the circuit was the strong negative temperature coefficient of $50\text{kppm}/^{\circ}\text{C}$ or greater than may be obtained by subtracting the output of a constant current reference from a reference with a negative temperature coefficient. Test circuits were fabricated, but operated at an incorrect current output level due to the inadvertent shorting of the P-MOS cascade transistor N-Wells. Despite this, the concept of temperature compensated bias voltage was demonstrated using a single PTAT current reference circuit combined with an off-chip reference to compensate for the large bias offset.

4.3 Demonstration System

We describe a new optoelectronic switching system demonstration that implements part of the distribution fabric for a large ATM switch. The system uses a single optoelectronic VLSI modulator-based switching chip with more than 4000 optical I/O. The optical system images the inputs from a two dimensional fiber bundle onto this chip. A new optomechanical design allows the system to be mounted in a standard electronic equipment frame. A large section of the switch was operated as a 208 Mb/s time multiplexed space switch, which is applicable to ATM switching using the appropriate out-of-band controller. A larger section with 896 input light beams and 256 output beams was operated at 160 Mb/s as a slowly reconfigurable space switch.

4.3.1 Introduction

In the past few years, the demand for telecommunications services beyond voice telephony has skyrocketed. For the growth of these services to continue at this rate, cost effective means of transporting and switching large amounts of information must be found. Although fiber optic transmission has significantly reduced the cost of transmission, switching high bandwidth signals remains expensive.

While all electronic switching systems are certainly possible for these high bandwidth systems, considerable effort has been expended to reduce the cost of fiber optic connections between frames or racks of equipment separated by several meters. As an example, one can envision fiber-optic data links connecting the line units that receive and transmit data from the outside world with an electronic switching fabric. Optical data links, ODLs, can perform the optical to electrical conversions. Several of these optical data links can be electrically connected with electronic switching chips on a printed circuit board.

As the demand for bandwidth increases, several hundred to several thousand optical fibers might be incident on the switching fabric. Discrete optical data links and parallel data links with up to 32 fibers per data link remain an expensive solution to transporting this information due to their per-link cost, physical size, and power dissipation. Power dissipation on the switching chips is high because of the need for electronic drivers for the high speed electrical interconnections between the switching chips and the data links. By integrating the O/E conversions directly onto the switching chips, lower cost and higher density systems can be built.

We demonstrated an experimental optoelectronic switching network based on this lower cost solution. This demonstration differs in many ways from our earlier system experiments. The system uses a new device technology consisting of GaAs/AlGaAs multiple quantum well modulators and detectors flip-chip bonded to silicon VLSI circuitry[100,19]. The system implements part of a new simplified distribution network for the growable packet architecture[101,102]. The distribution fabric is implemented with a single chip, contrasting with previous systems consisting of cascaded chips. The mechanical design of the system uses a plate-pedestal system [103], that provides superior robustness compared to the slot-plate systems [104]. This system is mounted in a standard electronic equipment frame. The system contains a single two dimensional fiber array providing fibers for the input signals and read beams and providing fibers for the output beams. The optical system images the inputs from the fiber bundle onto the switching chip, provides optical fan-out of the signals from the fibers to the switching chip, and images the outputs from the chip onto the fiber bundle. A large section of the switch was operated as a 208 Mb/s time multi-

plexed space switch, which is applicable to ATM switching using the appropriate out-of-band controller. A larger section with 896 input light beams and 256 output beams was operated as a slowly reconfigurable space switch.

The architecture is first described in the section that follows. An architecture can be judged (or compared against other architectures) by the ratio of the throughput to the cost. In packet switching, the throughput can be defined as either the bandwidth or the number of packets per second. If the packets have a fixed length, such as the case with ATM switching, then optimizing the number of packets per second simultaneously optimizes the bandwidth. A key in determining the "right" architecture is to make the correct tradeoff between cost and performance. Lets consider the simple case of connecting multiple computers. On one end of the spectrum, one can use a shared bus (such as Ethernet) in which case the performance is limited by the bandwidth of the shared medium. On the other end, one can provide a dedicated resources, both a high bandwidth connection and a dedicated memory buffer between all sets of two computers for optimum performance. The architecture on which our photonic switching demonstration is based, uses several novel ideas to achieve performance nearly equal to the latter case, but at a fraction of the cost. While it is beyond the scope of this report to fully describe the architecture, the highlights are given in the following section.

4.3.2 Architecture

Asynchronous Transfer Mode (ATM) is the leading approach to routing high bandwidth signals for the telecommunications networks of the future. As demands for wide bandwidth services grows, there will be a need for telecommunications switching networks with aggregate capacities beyond 1 Tb/s. While there are many approaches to a large capacity ATM network, the

architecture that we describe here is well suited to this task.




Figure 21. Growable Packet Architecture for a 256 input 256 output switching network

ATM switching networks must have memory for storing ATM cells in the event that two or more cells are destined for the same output at the same time. The size of the memory(s) is determined by the traffic statistics and enough memory must be provided so that the chances of an ATM cell being dropped should be very small, perhaps 10^{-9} . Typically, greater than 10000 cells might need to be stored per output. A simple method of building an ATM switch is to use a single large memory, which may contain many memory chips. The incoming ATM cells are sequentially written into the memory and the outgoing cells are sequentially read from the memory in the appropriate order, depending on their destination. This approach works for fairly modest sizes, but as the number of input and output ports becomes large, the memory access time becomes small. For example, a 2.5 Gb/s ATM switch with 256 input and 256 output ports requires an access time of less than 350 ps even if we could devise a way of writing entire ATM cells into and out of the memory in parallel. One way of partitioning the problem, so that one needs only modest speed memories is the growable packet architecture [101] shown in Fig. 21 for a 256 x 256 ATM network. The growable packet architecture consists of two subsystems. The first is a distribution network that provides fan-out and distributes the input signals to the second stage which consists of output packet modules. The distribution network contains no memory, so it consists of a non-blocking (or very low blocking) interconnection network. The output packet modules have buffering or memory as described above so that ATM cells destined for the same output port can be

stored temporarily and ATM cells are not lost.




Figure 22. Simplified Growable Packet Architecture for a 256 input 256 output switching network

The distribution portion of the network can be a challenge if the network becomes large, such as the one in Fig. 21. This is because of the large number of crosspoints and connections between them as well as the large number of calculations that need to be performed to route data through the network during the ATM cell period. The architecture that we have chosen for this demonstration uses a novel network that greatly simplifies the distribution network [102]. In Fig. 22, we show the implementation of the simplified 1024 x 1024 distribution network, consisting of four groups called pipes, each consisting of 16 16 x 16 switches. This network contains 1/16 of the number of crosspoints as a 1024 x 1024 crossbar, yet has very low blocking probability. The reasons for this are that the inputs to the individual switching chips are arranged so that two inputs incident on the same 16 x 16 switch in one pipe are not incident on the same switch in the other pipes, and that the routing algorithm ensures an even distribution of calls through the four pipes [105].

In our implementation of this network, the distribution network is designed to be controlled by an *out of band controller* [105]. In this type of network, the routing information from all the data inputs are routed to a *path hunt processor* that reads the header information and calculates the appropriate paths through the distribution network. It must be able to perform these calculations within the ATM cell time (173 ns for 2.5 Gb/s data, 2.5 μ s for 155 Mb/s data). The path hunt processor is designed to calculate the paths through the network using global information from all the inputs and it performs these calculations during the previous ATM cell period. This is in contrast to a self routing switch (in band control) that uses local information to route the data through the network. An out of band controller with this distribution fabric allows for simple multicasting

(sending one input to many outputs), fault avoidance, variable length packets, cell loss priority handling, and dynamic load balancing. A large advantage of this reduced crosspoint distribution fabric is that it allows the calculations of the paths through the network to be calculated in parallel [105], so that the path hunt processor, even for a 1024 input 2.5 Gb/s switching network, can be implemented with logic operating below 50 Mb/s.

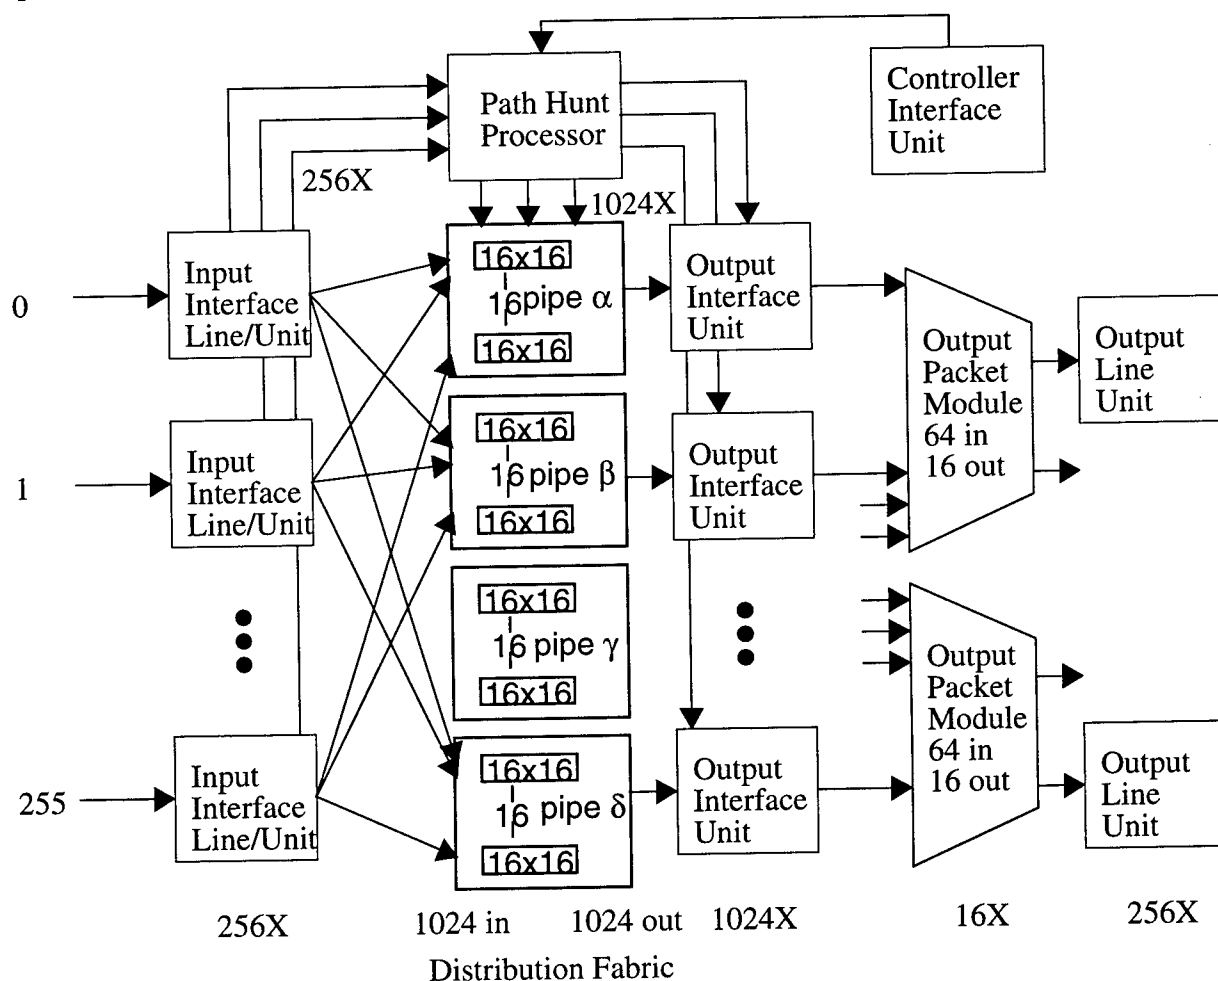


Figure 23. Switching system (vaporware) consisting of the distribution fabric, input and output interface units, path hunt processor, and output packet modules

A full system is shown in Fig. 23. Items in grey were not built. The input interface unit contains several functions. First, if the incoming data is from a SONET (synchronous optical network) link, the line unit must provide clock recovery, error detection, SONET pointer processing, and frame delineation. It would also extract the ATM cell, change the routing information contained in the header (VPI/VCI addresses) to a form that is relevant for the path hunt processor, perform a translation of the VPI/VCI addresses, and provide a small amount of buffering of ATM cells so that they can be stored temporarily during the path hunt processing. Our particular implementation also required the input interface unit to insert a guard band in the data during which the switch reconfigures, encode the data so that long strings of ones or zeros do not occur, add parity check bits, and add a preamble for synchronization at the receiving end. The 424 bit (53 byte) ATM cell at 155 Mb/s would be transformed to a 576 bit (72 byte) word at 208 Mb/s. The input

data lasers would also be driven from this unit.

The output interface units between the distribution unit and the output packet modules would then decode the inputs, re-synchronize the inputs and remove the guard band. It would also have to do a second translation of the ATM routing information. An output line unit, after the output packet modules would put the data back into the SONET format. A controller interface unit would provide an interface to a PC or workstation that would allow us to program which physical paths through the distribution fabric corresponded to which VPI/VCI values, as well as perform error checking and other diagnostic functions.

For our demonstration, we have chosen to implement one pipe of the distribution network, consisting of 16 16 x 16 switches, operating at the OC-3c rate of 155 Mb/s. A digital word generator, controlled by a personal computer, supplies the out of band control signals that the path hunt processor would supply in a completed system. It can supply time multiplexed data inputs formatted as the 72 byte cells that the input interface unit would have provided. The word generator can reconfigure the switching fabric between the cells. Thus, the experimental system demonstrates that ATM cells could be sent through the switch, assuming that no difficulties were encountered in building the appropriate input and output interface circuitry. Alternately, the switch can be configured as a space switch, with input data coming from any source, including digitized video as we will discuss in the last section.

4.3.3 Switching Chip

All 16 switches for one pipe of the distribution fabric are implemented on one optoelec-

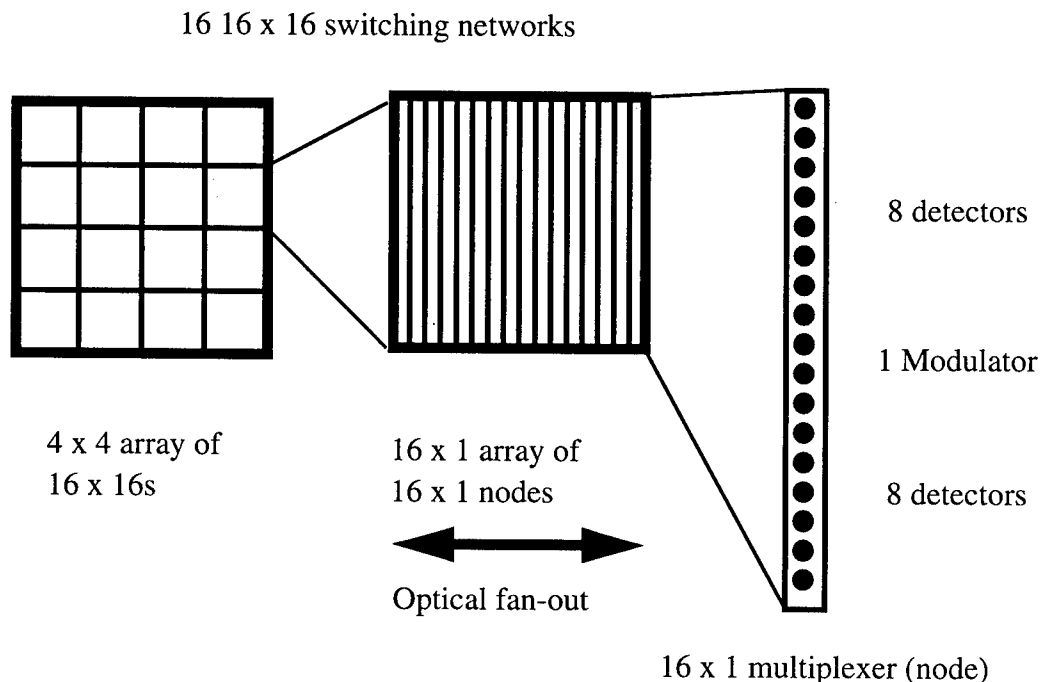


Figure 24. Schematic diagram of the Optoelectronic VLSI switching chip.

tronic VLSI chip shown in block diagram form in Fig. 24 [73]. The 16×16 networks are implemented with a passive optical fan-out and electronic fan-in using 16×1 multiplexers or switching nodes. The control of the switching chip is electronic. The switching chip contains $16 \times 16 \times 16 = 4096$ optical inputs and $16 \times 16 = 256$ optical outputs. An idle channel can be routed to any of the outputs, this idle channel is connected to the chip electrically. The control of the individual 16×1 nodes (17×1 if you include the idle channel) is provided by providing 17 primary memories and 17 shadow memories, one for each of the data inputs and one for the idle channel. Data is loaded into the shadow memories one column at a time (64 columns) and the data is transferred from the shadow to the primary memories during the guard band placed in the data by the input interface unit. There are 23 moderate speed (26 Mb/s) electrical inputs to the chip. Twenty of these contain the decoded electrical control information. The other three are the shift register clock, shift register input, and the idle channel.

The chip was designed using standard $1.0 \mu\text{m}$ CMOS. The 4352 optical detector/modulators were flip-chip bonded to the silicon CMOS circuit and the substrate removed to allow access the 850 nm optical ports. Details of the construction are given in reference [19]. The center to center spacing of the optical I/O was $80 \mu\text{m}$, the optical window size was $\sim 11 \mu\text{m}$, and the optical field of view was 5.44 mm . While several circuits have now been made using this process, this circuit has the most optical I/O and greatest electrical complexity at the time it was built. The chip is packaged on a custom aluminum mount with a four layer flexible microstrip circuit providing the control signal and bias connections. Bypass capacitors are connected between the V_{dd} , V_{P} , ECL(3V), $V_{\text{modulator}}$ and V_{detector} power supplies and ground, and 50 ohm resistors are connected between the input control signals and $V_{\text{P-ECL}}$. Detailed description and characterization of the chip is given in reference 73.

The chip implements part of an optoelectronic distribution network for an ATM switching demonstration. The circuit has 4096 optical detectors and 256 optical modulators. The control information is brought into the chip via electrical connections. It is decoded and routed to the individual nodes in a column by column basis using the parallel outputs from a shift registers.

Fig. 25 shows a schematic of a 16×1 switching node consisting of 16 receiver/selectors, an OR tree, control memories, and an output section. Throughout the rest of this section we will use the term "switching node" rather than the term "multiplexer" because a node contains elements, such receivers, control memories and modulator drivers, not normally associated with a multiplexer. The receiver/selectors serve two functions. First, they convert the 16 optical inputs into electrical signals. Second, based on information stored locally in control memories, they "select" which of the 16 inputs is to be routed to the output. That is, only one of the receiver/selectors will be enabled at a time. The outputs from the receiver/selectors are then routed to a 16 input OR gate tree, implemented with 4 stages of two input NAND/NOR logic and then routed to the modulator driver section. Using a fan-in of two in these gates minimizes capacitive loading because the gates are spread out in space across the node. The longest electrical trace is approximately $320 \mu\text{m}$, with an estimated capacitance of 25 fF . In the output section, the data passes through a 2×1 multiplexer and then to a final modulator-driver inverter. If none of the receivers are selected, the 2×1 multiplexer inserts an idle signal, which is routed onto the chip electrically.

Each individual optical input of the node has a shadow and a primary memory associated with it that determine if that particular optical input is the one of sixteen that is to be routed to the

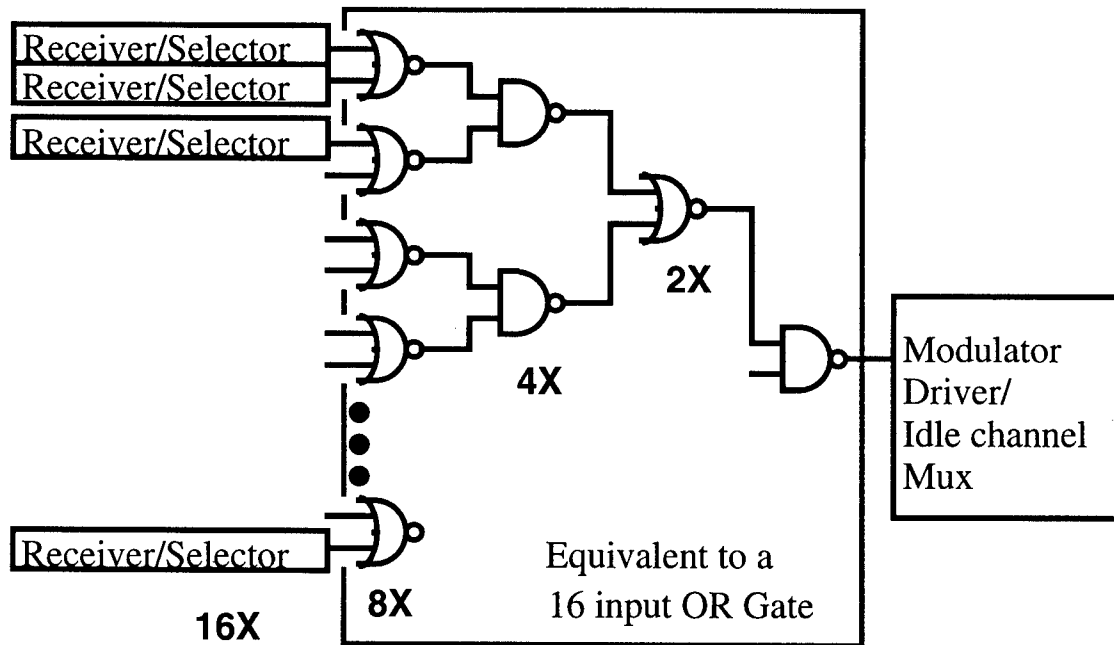


Figure 25. Schematic diagrams of a 16 x 1 node.




Figure 26. a) Receiver/selector and modulator driver/idle channel multiplexer schematic

node output as shown in Fig. 26. There is also a shadow and primary memory associated with the output section that determines whether the idle channel is selected as shown. The shadow memories are loaded while data is passing through the nodes and then the control information is transferred from the shadow to the primary memories during a guard band in the data.

The control information is read into the shadow memories in a column by column basis as illustrated in Fig. 24. Four 5:17 decoders on the left side of the array provide the control information which is routed horizontally to the nodes in each of the 4 rows. The 5 input bits per row provide control information for the 16 receiver/selectors plus one bit for the idle channel. A shift register provides an input to each column of nodes that enables writing of the shadow memories

with the control information bits, one column of nodes at a time. The 65th bit of the shift register provides a signal that transfers data from the shadow to the primary memories. Thus, the signal input to the shift register is a logic one followed by a string of logic zeros. The input data and control load signals are synchronized in time by a common control load clock and master-slave flip-flops in the control information inputs. Clock frequencies above 100 Mb/s have been used to load the control information into the switching nodes. This corresponds to a reconfiguration time of ~ 655 ns to load control information into the shadow memories in each of the 64 columns of nodes in the array and transfer this control information from the shadow to primary memories. The time required for the receivers to become active after this transfer step was measured to be less than 5 ns.

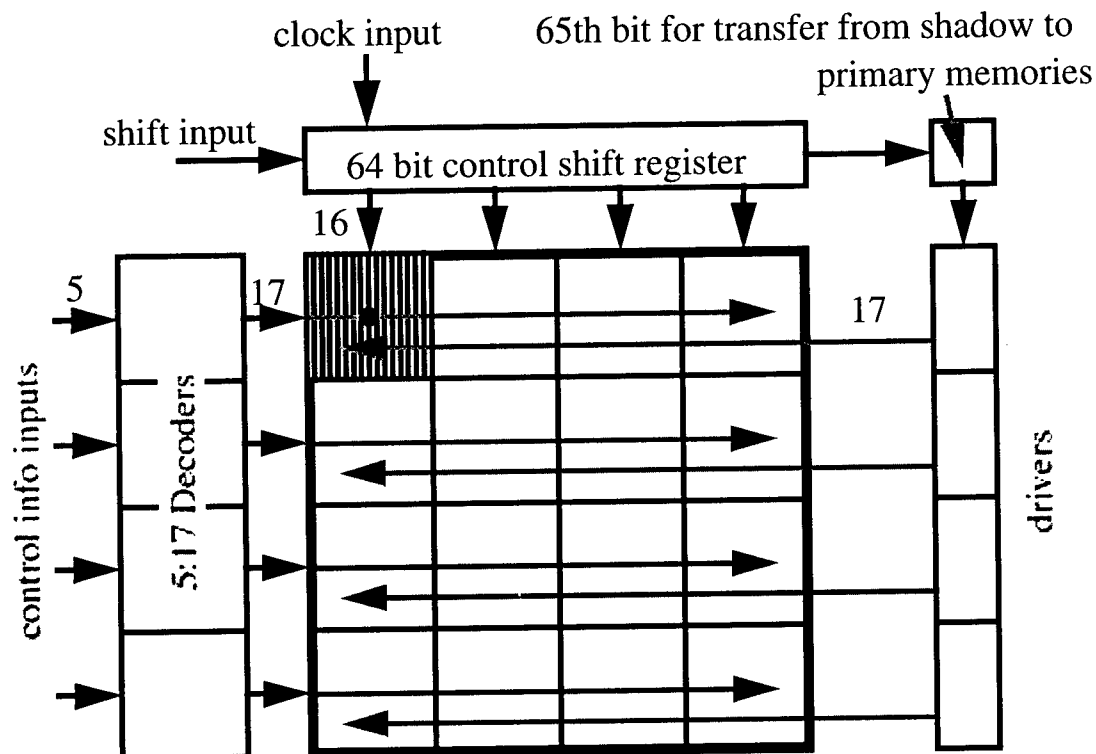


Figure 27. Block diagram of the electronic control circuitry of the chip. All control inputs are electronic. The control information is sequentially loaded into the shadow memories, one column at a time, using the outputs from each bit of the shift register to enable writing of the memories. A signal to transfer the data from the shadow to primary memories is derived from the 65th bit of the shift register.

Each electrical control bit is connected to one of two inputs to an electrical differential amplifier. A reference voltage is connected to the other side of the amplifier. The amplifiers unique design [110] enabled the threshold to be set anywhere from 0.5 to 4.5 volts (with a 5V supply), whereas standard differential amplifiers often must have a reference near the center of the voltage range. For compatibility with standard electronic circuitry, the reference voltage can be nominally set to 3.7 V, the threshold of positive emitter coupled logic (P-ECL). However, almost all of the testing was done with the reference voltage at 0.5V, with the input voltage swing from 0 to 1V.

Each of the 4096 optical receivers is a DC coupled transimpedance design with a novel non-linear feedback element to improve the dynamic range [109]. This feedback element consists of a parallel combination of a p-type FET with its source grounded that acts as a resistor for small photocurrents and an n-type FET with its source connected to its drain that acts as a voltage clamp that limits the voltage swing for larger photocurrents. Because only one of 16 inputs is routed to the node output, the receiver resembles a two-input NAND gate. One input to the NAND gate is the detected photocurrent and the other input is the signal from the control memory, that determines, based upon the control information, whether that input is the selected one. Performing the selection process in this way has the important advantage of reducing the static dissipation in the unselected receivers. Because each active receiver dissipates ~ 2.5 mW, this reduces the static dissipation of the chip from ~ 10 W if all receivers were continuously biased to less than 1 W when only 256 are selected. The dissipation was determined from SPICE simulations.

The metallic pads for bump bonding are $15\text{ }\mu\text{m} \times 15\text{ }\mu\text{m}$ with a $15\text{ }\mu\text{m}$ space between the n-type and p-type connections of a diode. The individual diodes are on $80\text{ }\mu\text{m}$ centers. The active area of the optical window is $11\text{ }\mu\text{m} \times 11\text{ }\mu\text{m}$, which is reduced from the pad and diode size by an isolation implant. No circuitry was placed underneath the bump-bond pads in this design, even though we have now made circuits with FETs underneath the pads [109]. In the vertical direction, the inputs to a particular node are arranged in a column of 17 diodes consisting of 8 detectors, the output modulator, and 8 detectors as shown in Fig. 27. Since there are 4 nodes in the vertical direction, there are 68 MQW diodes down a column. In the horizontal direction, there are 64 nodes, but they are arranged in four groups of 16 with a gap of $80\text{ }\mu\text{m}$ between groups. The gap provides space for V_{dd} and ground power supply connections, so that the array can be powered in sections and voltage variations on the power supply leads are minimized. The optical field of view is $67 \times 80\text{ }\mu\text{m}$ or 5.36 mm in the horizontal direction and $68 \times 80\text{ }\mu\text{m}$ or 5.44 mm in the vertical direction. The shift register, decoder, transfer lead drivers, test circuits, and electronic I/O circuitry surround the smart pixel elements of the array. The total chip size is $7\text{ mm} \times 7\text{ mm}$.

The chip is packaged on a custom aluminum mount with a four layer flexible microstrip circuit providing the control signal and bias connections. Bypass capacitors are connected between the V_{dd} , $V_{P-ECL}(3\text{ V})$, $V_{\text{modulator}}$, and V_{detector} power supplies and ground, and $50\text{ }\Omega$ resistors are connected between the input control signals and ground. To interface with standard test equipment, we chose to connect the resistors to ground rather than to V_{P-ECL} .

All but two of the optical modulators and detectors generated photocurrent in response to incident light. The voltage dependence of the reflectivity and responsivity of the modulators and detectors can be measured by sweeping the voltage between the detector power supply and ground. Although there is no direct connection of the photodiodes to ground, the receiver circuitry completes the connection. First, a forward biased diode in series with the photodiodes, from the parasitic diodes between the p-diffusion and n-well of the p-type feedback FET, that provides a current path to V_{dd} . A non-linear "resistor", from the static current versus voltage characteristics of the receivers and electrical differential amplifiers, completes the connection between V_{dd} and ground. The reflectivities in Fig. 28 are flat for voltages below 2.5 V , because the 0.7 V drop across the forward biased diode and 1.8 V drop from V_{dd} to ground reduce the voltage that appears across the photodiode from that supplied to the circuit. Nonetheless, the data can be used to compare the uniformity of the detectors and to measure high and low state reflectivities.




Figure 28. Reflectivity versus voltage for the MQW diodes on the four corners of the array, measured at ~ 851.5 nm and ~ 20 μ W.

The reflectivities for the detectors on the four corners of the array are shown in Fig. 28. The thickness of the antireflection coating was not optimum, limiting the contrast ratio to less than 2:1. The high and low state reflectivities at 6 and 11 V were 0.44 ± 0.03 and 0.25 ± 0.03 . The uniformity in reflectivities is much better than in our previous devices[93], because of better uniformity of the thickness of the stop etch layer. If the A/R coating is not perfect, Fabry-Perot resonances will be present, and variations in cavity length will vary the resonant frequency, which in turn varies the reflectivity at a fixed wavelength.

In Fig. 29, we show the normalized output from one 16×1 node from each of the 16×16 sections in the array, with 2 of the 16 inputs active at a data rate of 200 Mb/s. Input 0 had a pattern of "11100010" and input 4 (0100) had a pattern of "01010011". In this figure, all devices had the correct output bit pattern. We also measured all 256 16×1 nodes with 4 of 16 inputs active at 200 Mb/s. One of the 20 decoded control bits was stuck in a fixed state, preventing selection of 8 of the inputs (2 of which were measured) from the nodes in the bottom row of the array. Other than that, all devices had easily recognizable bit patterns. There was significant output amplitude variation from device to device, caused by defocus and positioning errors as the motorized stages moved the array across the fixed input and read beams. In particular, the defocus caused the amount of light coupled into the output fiber based detector to be reduced. The data in Fig. 28 indicates this variation is not present in the devices, because the reflectivities of the modulators are fairly uniform and they are driven by voltages that should not vary in amplitude. Normalizing

the data in Fig. 29 caused the “noise” to be magnified toward the lower right corner where the amplitude was reduced. The delay variation in these measurements of ~ 1 ns was also likely caused by positioning errors leading to variations in photocurrent. More detailed measurements on uniformity and delay variations and are given below.




Figure 29. Normalized output bit patterns at 200 Mb/s from one 16 x 1 node from each of the different 16 x 16 switches (see Fig. 22). Input 0 had a pattern of “11100010” and input 4 had a pattern of “01010011”.

In Fig.30, we show a superposition of 16 eye diagrams at 400 Mb/s of a 16:1 node with one input selected at a time. The photocurrent was monitored as the input spot was moved from detector to detector to ensure positioning errors did not contribute to eye closure. The eye depicts the combined jitter, skew, and pulse width distortion for an entire node and has sufficient opening for reliable operation. One could achieve a clean eye diagram at data rates up to ~ 470 Mb/s. The limiting factor is the modulator driver which was an inverter with $3\text{ }\mu\text{m}$ wide FETs. There was not additional space to make a larger driver.

Next, we looked at the dependence of the pulse width on the node bias voltage (V_{dd}) and on optical power. Pulse width distortion places a limit on the maximum bit rate that can be achieved in the system application of these arrays. The measured pulse widths were different for the devices in the even and odd columns of the array, because two different feedback resistors were used. This was accomplished by varying the gate length of the p-type FET from $1\text{ }\mu\text{m}$ to $1.5\text{ }\mu\text{m}$. The feedback resistor determines the optical input power or current that causes the output of the receiver to change from a low to high value (i. e. the receiver threshold). The circuit was

designed with two different feedback resistor values in case the threshold ended up too high and we did not have enough optical power or in case the threshold ended up too low and the RC time constant was too long. This was unnecessary, because the threshold can be adjusted by varying V_{dd} . The dependence on V_{dd} occurs because the threshold of the receiver NAND gate is a function of V_{dd} , so the nominal gate to source voltage of the p-type FET changes as a function of V_{dd} and thus its effective resistance changes. Higher values of V_{dd} should increase the gate to source voltage and thus lower the feedback resistance, thereby raising the effective optical power threshold of the gate. Thus for a given delay or pulse width, higher optical powers are needed for higher V_{dd} . This was confirmed experimentally.




Figure 31. Solid lines show pulse width versus optical input power at various values of V_{dd} from 5.0V to 6.0V for a node with the smaller effective feedback resistance value near the upper left corner of the array. Solid circles indicate the actual data points. Power was assumed to be twice the photocurrent, which was monitored during the set of measurements. The dotted lines, with point labels +, o, and x, correspond to nodes near the other three corners of the array at $V_{dd} = 5.4$ V. The time resolution of the measurements was 198 ps

We modulated the input lasers at 200 Mb/s with a pattern consisting of "0000100011110111". This pattern gives a lone "1" (the 5th bit) and a lone "0" (the 13th bit). Looking at the width of these bits gives a good indication of pulse width distortion. Fig. 31 shows the pulse width of a node near the upper left corner of the array with one particular input selected for the 5th bit in the pattern (the lone "1") for various values of V_{dd} versus optical input power.




Figure 30. Sixteen eye diagrams at 400 Mb/s *superimposed*, where each eye diagram is the optical output from a 16 x 1 switching node with one of its 16 optical inputs illuminated with pseudorandom data with a word length of 2^{23} .

The time resolution of the data was 195 ps. One would expect that if the lone “1” bit had a longer pulse width, the lone “0” bit would have a shorter pulse width and that the average of the two would be 5 ns. Indeed the average of the two pulse widths was 5 ns to within 310 ps. The 3 dotted lines in the figure show the same data for nodes near the other three corners of the array, measured over a smaller power range. The total variation of the four nodes for a given optical power and voltage was less than ± 400 ps. The variation can be caused by differences in the transistor characteristics across the array or by variations in V_{dd} or ground potential across the array. Because the static current is low, it is likely the former. It is unlikely that random variations exist across the array, because the transistor characteristics tend to vary in a smooth fashion for an established CMOS process.

The data in Fig. 31 also shows that the allowed variations in optical power is greater for higher V_{dd} values or higher thresholds, even when plotted on a logarithmic scale. This occurs because the non-linearity in the feedback resistor “compresses” the voltage swing as a function of power. For example, at high powers, a 20% increase in optical power might only cause a 10% increase in voltage swing, whereas, at lower powers, this same increase in optical power would cause a 20% increase in voltage swing. Because the delay and pulse widths are functions of the voltage swing, a reduction in voltage swing versus optical power translates into less variation in

pulse width as a function of optical power.

Element	Cap (pF)	Rate (Mb/s)	transitions per bit	diss./line (mW)	number	diss total (mW)
Control load	3	25	0.031	0.029	65	1.88
transfer bit	3	25	0.031	0.029	68	1.96
decoder bits	0.8	25	1	0.250	10	2.50
control information	3.7	25	1	1.156	68	78.63
input pad (dynamic)	2	25	1	0.625	23	14.38
idle channel	2.5	50	1	1.563	4	6.25
clock	3	25	2	1.875	2	3.75
rcvr-on (dynamic)	0.12	200	1	0.300	256	76.80
modulator(dynamic)	0.06	200	1	0.150	256	38.40
rest of node	0.25	200	1	0.625	256	160.00
rcvr-on (static)		200		2.500	256	640.00
rcvr-off (static)		200		0.060	3840	230.40
modulators (static)		200		0.400	256	102.40
input pads(static)		25		6.000	45	270.00
Total dissipation						1627.34

Table 12. Calculated chip dissipation for various sections of the chip, assuming the entire chip had optical inputs and was operating. The dynamic dissipation is equal to $1/2NCV^2B_{eff}$, where C is the capacitance, V is the voltage swing (5V), B_{eff} is the effective bit-rate, which is equal to the actual bit-rate (200 Mb/s for data, 50 Mb/s for idle channel, and 25 Mb/s for control), times the number of transitions per bit (for the control loads, we assume that there is one pulse or 2 transitions every 65 bits, for RZ data there are 2 transitions per bit, for NRZ data there is 1 transition per bit), and N is the number of lines of the chip. The Capacitance was estimated by summing the total gate area capacitance and twice the extracted values (to be conservative). Static power dissipations were simulated using SPICE. (The 45 input electrical pads include 22 test pads) We do not know the exact dissipation of the disabled receivers that were not illuminated.

The measured static power dissipation of the array ranges from 750 mW at $V_{dd}=5V$ to 1.5W at $V_{dd}=6V$. Roughly 70% of this is in the receivers and 30% is in the input differential amplifiers for the electrical control signals. The dynamic dissipation can be estimated by $1/2 CV^2B_{eff}$ where B_{eff} is the effective bit rate of the particular signals. In Table 12, we show the design data rates for the various parts of the circuits and the calculated dynamic dissipations and compare them to the static dissipations. The static dissipations were simulated using SPICE. The unselected receiver dissipation, assuming that input light is present, is equal to the product of the photocurrent and the difference between the detector voltage and V_{dd} , because the parasitic diodes of the feedback FETs provide a path for the photocurrent to V_{dd} . For an average photocurrent of 20 μA and a voltage difference of 3.0V, the 3840 unselected receivers have a dissipation of 230 mW. The static dissipation from the 256 optical modulators, for a photocurrent of 50 μA with a voltage of 8.0V, is

102 mW. From the table, one can see that the ~2.5 mW static power dissipation of the receivers dominates. This does not mean that optical interconnections are not warranted, indeed, the use of optical interconnections greatly reduces dynamic dissipation by eliminating the need for long electrical traces on the chip (and large transistors to drive them) as well as large dissipation hungry electronic output drivers. It may be more optimum from a power dissipation point of view to have more electronics per optical I/O. Others have also reached this conclusion[111].

Incidentally, during testing, only one or two nodes had optical inputs. Without optical inputs, the static dissipation of the unselected receivers and modulators and the dynamic dissipation of the switching nodes (including the receivers and modulators) are both approximately zero. For this reason, it is important to build systems to test the entire array concurrently to demonstrate that power dissipation will not be a problem.

Two cases of crosstalk were measured using two lasers with slightly different bit rates incident on the devices. In one case, the interfering signal was incident on detectors within the same node, and in the other case the interfering laser was incident on detectors in a neighboring node. With the pseudo-random optical data inputs at 200 Mb/s, no eye closure is observed in either case. However, by looking at the detected optical output on a spectrum analyzer, crosstalk, 45 dB below the signal, could be observed at 200 MHz for square wave inputs when the interfering signal was incident on a selected receiver of a neighboring node. No observable crosstalk was seen when the signal was incident on an unselected receiver, either in the same node or a neighboring node.

Voltage variations due to simultaneous switching currents through the parasitic inductances and resistances of the supply lines is the most likely reason for the observed crosstalk. The modulator driver supply lines are more likely to contribute crosstalk than as the receiver supply lines, because the crosstalk was independent of the selected receiver on the neighboring node. If we extrapolate the crosstalk value measured and assume each of the 16 nodes in a row on a common bias lead will contribute the measured amount, the overall signal to noise ratio should be $-45 \text{ dB} + 10 \log(16) = -33 \text{ dB}$. This would cause only a 0.20 dB power penalty for an input noise limited receiver with an incident signal that comes from the node output.

4.3.4 Switching Chip Mounts

For some optoelectronic devices, the optical characteristics of the device are a strong function of the device temperature. An example of this is the self-electro-optic effect device (SEED)[112]. SEEDs make use of the shift in wavelength of the exciton absorption maxima that occurs as a function of a changing electrical field across multiple quantum well material[113]. For a typical device, the absorption maxima must be shifted by 3 to 5 nm to obtain good contrast between the absorptive and reflective states. A change in the device temperature, however, can also change the location of the exciton absorption maxima[114]. For GaAs/AlGaAs devices the absorption maxima shifts approximately $0.28 \text{ nm}/^\circ\text{C}$ [115]. Thus, it becomes necessary to carefully design the SEED mount to minimize the temperature gradient and, therefore, this temperature induced shift across the chip.

In this paper we have used finite element analysis (FEA) to model mounts for a 16x16 array of FET-SEED switching nodes. By careful mount design, the calculated temperature spread could be held to 1°C even when the power density was $40 \text{ W}/\text{cm}^2$ over the 0.15 cm^2 active chip

area; a 6W chip. We have also used the temperature dependence of the exciton absorption maxima to map the temperature of an existing 4x4 array of FET-SEED switching nodes [116], operating at a power density of $49\text{W}/\text{cm}^2$ and found the results in good agreement with those obtained by FEA.

Two fundamental tasks exist in the thermal management of a chip; not only is it necessary to prevent the entire chip from heating to a point where thermal effects degrade the overall performance but it is also necessary to maintain the temperature of multiple locations on the chip to nearly the same temperature. The exact amount of temperature variation that can be tolerated (ΔT) will depend on the application but in general it will fall into the range of 1 to 4°C . It is not necessary to hold the entire chip to this temperature range but only specific devices on the chip, such as all the optical output modulators. The temperature variation of interest is the difference between the hottest and coolest of these devices, as shown in Figure 32, not the overall temperature spread or the temperature differential of a single node. The temperature profile depicted in Figure 32 is




Figure 32. Schematic 1D Temperature Profile SEED Node Array.

representative of that which one would expect in a regular array of nodes. The chief cause of temperature variation between equivalent devices on the chip is the spreading of heat to the predominately passive regions around the periphery of the chip where the electrical I/O bond pads are located.

The design of a SEED mount must, therefore, contain an analysis of the temperature variation between the equivalent critical optoelectronic components on the chip. The testing of the mounted chip must contain an analysis to verify the temperature variation between these components. The first of these two tasks can be accomplished by finite element analysis and will be discussed later. The second of these two tasks can be accomplished using the same physical phenomena that makes attention to temperature variation necessary, namely the shifting of the exciton peak location. The experimental setup used in this study is shown in Figure 33. The chip, a 4x4 array of $210\mu\text{m} \times 210\mu\text{m}$ FET-SEED switching nodes [116], was held at a constant temperature with a thermal electric cooler (TEC). A laser light source, $\lambda=850\text{nm}$, and lens system was used to illuminate a single SEED modulator. The reflected light was focused on a photodetector. A bias voltage ramp was applied to the modulator and the photocurrent of the reflected light measured as a function of the applied bias. The minimum photocurrent corresponds to the absorption maxima. Next, the temperature of the entire chip was changed by adjusting the TEC. The change in bias voltage at which the absorption maxima occurred, for the same modulator, was noted. In this manner a ΔV versus ΔT curve was constructed. Finally an X-Y stepper stage was used to move the laser illumination to different modulators and the bias voltage at which the absorption maxima occurred noted as a function of chip location. A finite element analysis (FEA) program was used to model the temperature profile of the measured device. The device mount is depicted




Figure 33. Experimental Setup

in Figure 34. The thermistor used to control the TEC, was not part of the thermal model but is




Figure 34. Thermal Model for FEA

included to further define the experimental setup. The thermistor was held at 20°C. The chip was powered to 350mW and the power was assumed to be distributed evenly over the 840um x 840um active area of the array (49W/cm²). The overall chip size was approximately 2.8mm square and the ceramic was 7.5mm square. Figure 35a is a plot of the measured temperature for four different modulators on the chip. Figure 35b is a plot of the isotherms on the active area of the chip surface




Figure 35. a) Experimentally Determined Temperature Profile 4x4 SEED Array

b) Finite Element Analysis Temperature Profile 4x4 SEED Array

as predicted by the FEA program. The thickness of the thermal epoxy used to affix the device to the ceramic was an estimate and could easily account for the minor differences between the model and the experimental measurements.

The chip just examined had a temperature variation across the active area on the order of 1.5°C, which for most applications would be acceptable. As mentioned, the chip was a 4x4 array of SEED nodes. For a chip of this size, even the center nodes are close to the non-active boarder regions and the effect of thermal spreading is therefore minimized. The temperature distribution

for a 16x16 array of 240um x 240um FET-SEED switching nodes, operating at 6.3W (40W/cm² over a 0.15cm² active area) mounted the same as the 4x4 array, is shown in Figure 36. Not only




Figure 36. Figure 5 Finite Element Analysis Temperature Profile 16 x 16 Node SEED Array

does the overall temperature of the chip rise but the temperature variation increases to 5°C. Again, the chief cause of the temperature variation is heat spreading into the inactive boarder regions of the chip; the active chip area is 3.84mm x 3.84mm and the overall chip size is 4.32mm x 4.32mm. A mount for the 16x16 node array was designed that would counter the effects of the heat spreading, as shown in Figure 37. An opening is cut in the ceramic used for the device interconnect and




Figure 37. Figure 6 Heat Confining Chip Mount

the device is mounted on a molybdenum pedestal that tapers down as it approaches the TEC. The taper in the pedestal counteracts the effect of the inactive chip periphery and the temperature spread is reduced to approximately 1°C, as shown in Figure 38. The decrease in the overall chip




Figure 38. Figure 7 Finite Element Analysis Temperature Profile. 16 x 16 Node SEED Array With Heat Confining Chip Mount

temperature is due to the use of solder for the two bonds in the thermal path instead of organic adhesives. The use of low temperature solders is required because of temperature limitations imposed by the TEC.

In conclusion, we have mapped the temperature across the active area of a SEED array by calibrating the exciton peak shift, of the GaAs-AlGaAs quantum well modulators, as function of temperature. We have compared the measured temperature profile to one modeled by finite element analysis and found them in good agreement. Using the finite element analysis program, we have shown that with proper design of the chip mount, areas of equivalent optoelectronic functions can be held at a uniform temperature, $\Delta T \sim 1^\circ\text{C}$, even for large high power density devices.

4.3.5 Optical System

In this section, the optical design of a photonic switching system is described in detail. The system is a 4-f imaging system with one lens design based on a Petzval lens and the second lens design based on a double Gauss lens. The assembled system is 60x70x370 mm and features a single imaging stage, pupil division for beam combination, and a single fiber bundle for both input and output. A brief description of the system requirements is provided, and the implication of some of these requirements on the paraxial design is explained. The choice of design forms for the lenses is explained, the design process is outlined, and the final lens designs are provided. Lens tolerancing is outlined and the final lens tolerances are provided. Finally, the assembly and alignment processes are described.

Introduction

In free-space photonic switching [14,104,117,118] data is encoded onto arrays of beams of light. These beam arrays are routed through the switching fabric and focused by lenses to form arrays of spots. The spots are registered onto arrays of detectors and modulators which perform switching functions. Arrays of optical fibers are used to input and output the data. This light data encoding and manipulation permits free-space photonic switches to perform the high-bandwidth interconnections that are projected to be needed in future telecommunication and data-communication switches.

Stringent requirements are placed on the optical systems in photonic switching systems. These requirements are necessary to ensure that the array of spots is imaged onto the array of detectors and modulators. The first requirement is that the images must be diffraction-limited; this requirement is necessary because the detectors and modulators can be as small as 5 μm . The second requirement is that the lenses maintain this diffraction-limited image quality over a large image field; this requirement is necessary because the complete array can be as large as 7 mm square. The third requirement is that the optical system must have well-controlled distortion and must provide precise alignment mechanisms; this requirement is necessary to achieve registration between the arrays of spots and the arrays of detectors and modulators. The fourth requirement is that the optical systems must be compact and robust; this requirement is necessary because photonic switching systems must work inside of electronics frames, a fairly small and hostile environment. Finally, the all of these requirements must be made in systems that meet cost constraints set by competing products.

The optical system requirements of a photonic switching system present many challenges to the optical engineer. Careful paraxial design is required to ensure that the optical system can perform the desired functions while ensuring that the system can be constructed. Diffraction-limited lenses must be designed to ensure that the focused light spots fall entirely within the active windows of the detectors and modulators. Methods for aligning the optical system must be

devised to achieve a high level of image alignment and to maintain that alignment over long periods. Assembly methods must be devised which allow precise construction and alignment while the optical system is in an electronics frame.

The many different groups working on photonic switching have developed many different methods for addressing these optical engineering challenges. Free-space interconnections such as the one described in this paper have been shown to be successful for large interconnections and for long distances.[119,120,121] Many other interconnection schemes show promise for certain applications. For interconnections between closely-spaced chips, interconnections directly through stacked chips[122] shows promise; planar optical interconnections[123] also have been demonstrated in this regime. For smaller numbers of interconnections, optical waveguides[124] have been demonstrated.

In this section we present optical design details of our most recent free-space photonic switching system. This optical design differs from our previous designs because it requires only a single imaging stage; the design of earlier photonic switching systems required as many as six cascaded imaging stages. Other important features of this system include: a single fiber bundle for both input and output optical fibers, an integrated array of diffractive microlenses to match numerical aperture of the fibers to the numerical aperture of the optical system, pupil division optics with offset lenses for beam combination, a multilevel diffraction grating and lens in the same substrate, and an overall mounted size of 60x70x370 mm.

The special features of this optical design played an important role in the successful demonstration of a photonic switching system that reliably routed 14 input signals at a data rate of 208Mbit/s.[14] This high performance was achieved while the system was mounted in a standard electronics frame in a room experiencing temperature fluctuations as large as 25 deg. F.

Functionality

The optical system described here is part of a growable packet switching architecture[125]. This architecture was developed to overcome some of the problems associated with high-capacity ATM switching systems. The architecture for this switch consists of three parts: a distribution network, a set of 16 output packet modules, and a network controller. In this case, the distribution network consists of four "pipes," each pipe consisting of sixteen 16x16 crossbars. The optical system described in this paper represents one of those pipes. A more complete description of the overall system is presented in another paper.[14]

Each input signal comes into the optical system on an optical fiber and is fanned out to 16 output channels, each on its own fiber. In this system, this functional fanout is performed by a 1x16 diffraction grating[126].

To provide a small spot in the object plane, the input signals are brought into the system on single-mode fibers. The optical fiber is arranged into a precision fiber bundle[107,127], which forms an array of spots in the object plane. Lenses image the array of spots onto an optoelectronic VLSI chip[73], which consists of an array of detectors and modulators on top of CMOS circuitry. The spots from the signal fibers illuminate detectors on the OE-VLSI chip. Light from additional fibers, called read fibers, follow a similar path through the optical system as the signal beams, but illuminate modulators on the OE-VLSI chip. The chip modulates the proper modulator, and light

from this modulator is reflected from the device. The lenses then re-image the reflected light from the OE-VLSI back onto an array of output fibers. The signal is brought out of the system on multi-mode fibers to allow for slight misalignments and image degradation. The output fibers are placed in the same fiber bundle as the input fibers to simplify the optical design and the fabrication of the fiber array.




Figure 39. Schematic showing the paraxial properties of the optical system. A/S 1 is the bottom of the aperture stop for lens 1, and A/S 2 is the bottom of the aperture stop for lens 2. In the paraxial design, the read beams follow the same path as the signal beams. Note that the optical axes of the lenses are offset to create separate paths for the input and output beams.

Figure 39 summarizes the paraxial properties of the system. Light enters the system on a single-mode fiber with a numerical aperture of 0.15, perpendicular to the object plane. The micro-lens array converts the numerical aperture of the beams to 0.04. The beams are collimated by lens 1, and pass through the center portion of the aperture stop for lens 1, which is in the same plane as the aperture stop of lens 2. The beams pass through half of the aperture stop of lens 2, which contains a diffraction grating. Lens 2 focuses the light onto the OE-VLSI, which reflects the light into the second half of the aperture stop of lens 2. This half of the aperture stop contains a weak lens. Lens 1 then focuses the light back onto the fiber face, and the light is coupled into a multi-mode fiber. The weak lens in the aperture stop is needed because there are no microlenses over the multi-mode output fibers.

Paraxial Properties

The first paraxial consideration in designing this system was to devise a way for the output beams to exit the system on different fibers from the read beams. In many previous systems, the output plane had to be physically separate from the input plane to provide input to additional stages of the switching system; some of these systems used polarization beam combining or amplitude division to accomplish this separation[117,118]. This system uses pupil division. The input signals and read beams pass through one part of the pupil; the output signals pass through a second part of the pupil. A 1x16 diffractive array generator in the input part of the pupil provides the functional fanout of the switching system.

The wavelength of the lasers was chosen to be 852nm. This wavelength falls within the bandwidth of good performance of the GaAs-AlGaAs multiple-quantum-well detectors and modulators, and coincides with the wavelength of high power, commercially-available, DBR laser diodes.

The paraxial properties of the object plane, the image plane, and the aperture stop are used as constraints in the lens design process. Figure 40 shows these constraints. The paraxial properties of the aperture stop are used as constraints in the design of both lens 1 and lens 2; these properties include the slope of the marginal ray, u_s , and the slope of the chief ray, \bar{u}_s . Additional constraints for the design of lens 1 include the paraxial properties of the object plane; these properties include the half object height, \bar{y}_o , and the numerical aperture, \bar{u}_o . Additional constraints for the design of lens 2 include the paraxial properties of the image plane; these properties include the half image height, \bar{y}_i , and the numerical aperture, \bar{u}_i .




Figure 40. The paraxial properties that are used for lens design constraints. \bar{y}_o and \bar{y}_i are the object and image heights, respectively. u_o and u_i are the numerical apertures in the object and image planes, respectively. \bar{u}_o and \bar{u}_i are the slopes of the chief ray in object and image space, respectively. u_s is the slope of the marginal ray in the aperture stop. \bar{u}_s is the field of view of the lenses.

The paraxial properties of the image plane are determined by the optoelectronic VLSI switching chip. The halfsize of the image is half the diagonal of the OE-VLSI chip, which is determined by the number and spacing of the detectors and modulators. In our system, the chip is a 64x64 square array of detectors and modulators on 0.080mm spacing, giving a half size image height of $\bar{y}_i = 3.62\text{mm}$. The $f/\#$ of the beam in the image plane is determined by the size of the detector windows and the wavelength of the lasers; in this system, an $f/4$ beam was chosen to ensure that the light from the 852nm lasers would fit well within the $10\text{ }\mu\text{m}$ windows. The location of the entrance pupil is determined by how the OE-VLSI is used; because the chip is used in reflection, the image plane must be telecentric. Finally, because the input beams use only half the aperture stop of lens 2, the $f/\#$ of the focusing lens must be half the $f/\#$ of the beams, so lens 2 must be $f/2$.

The paraxial properties of the object plane are determined by the properties of the fiber bundle. The half size of the object is half the diagonal of the fiber bundle, which is determined by the number of fibers and the spacing of the fibers. In this system the fiber bundle was designed to

be a 64x64 square array of fibers on a 0.250 mm spacing, giving an object size of $\bar{y}_o = 11.3$ mm. The $f/\#$ of the beams in the object plane is determined by the required $f/\#$ of the beams in image space and the system magnification; in this system, the beams in object space are $f/12.5$. A micro-lens placed over the fiber bundle face converts the $f/\#$ of the beams coming out of the fibers to the required $f/\#$. The object plane is telecentric because the fibers are perpendicular to the fiber bundle face. The collimating lens must be designed with an $f/\#$ 1/3 that of the beams leaving the fiber bundle because only 1/3 the diameter of the aperture stop is used by the beams in the input, so lens 1 must be $f/4.2$.

The paraxial properties of the aperture stop are determined by several factors. To simplify alignment and testing of the system, the aperture stop is placed in collimated space. The slope of the chief ray is determined by choosing a reasonable field of view for designing the lenses; semi field of view of $\bar{u}_s = 6.4$ deg. was chosen as a reasonable design task. This field of view, for the fixed object and image heights, forces the focal length of lens 2 to be 34mm and the focal length of lens 1 to be 106.25mm. Given the $f/\#$ of lens 1 and lens 2, the entrance pupil diameter of lens 2 becomes 17mm and the entrance pupil diameter of lens 1 becomes 25.5mm.

Lens design

To image the fiber array onto the detector array two objective lenses were designed. In addition to the paraxial constraints described in the previous section, the presence of the 1x16 diffraction grating forces several constraints on the lens design. First, $f \sin(\theta)$ mapping is required to ensure that each diffraction order lies directly on a detector or modulator. Second, the lenses must be individually well-corrected because the fan-out of the grating prevents aberration balancing across the field of view. Third, both of these individually-corrected lenses must have an external aperture stop, so that the diffraction grating can be in the aperture stop of both lenses. Finally, the aperture stop distances of the lenses must be large to allow space for elements needed for alignment, such as an afocal pair, a pair of wedges, and a pellicle beamsplitter; the large aperture stop distance also allows the two lenses to be readily combined into a single system.




Figure 41. Schematic of the 34mm lens. The long aperture stop distance is particularly helpful when aligning the spots onto the detectors and modulators because it allows a pellicle beamsplitter to be inserted for a viewport.

Table 13. System prescription for the 34mm lens. SFL6 is a high-index glass from Schott. Although it is relatively expensive at 6.7 times the price of BK7, its high index ($n_d = 1.805$) and relative ease of grinding and polishing make it an attractive glass for high-performance lenses such as this one.

surface	radius	thickness	glass
stop	INFINITY	0	
2	INFINITY	43.203	
3	51.191	7.000	SFL6
4	-227.808	4.500	
5	-34.201	20.000	BK7
6	-39.525	8.879	
7	29.658	20.000	SFL6
8	INFINITY	5.698	
9	-29.658	1.250	SFL6
10	54.033	2.250	
image	INFINITY	0	




Figure 42. Wavefront aberration plots for the 34mm lens at 852nm, for field angles of 0, 4, and 6.5 deg. Comatic aberrations are evident full field, and third-order aberrations are balanced by higher order terms.




Figure 43. The assembled 34mm lens. To simplify mounting hardware, the outer diameter of the barrel was designed to be the same as the diameter of the 106.25mm lens barrel.

34mm lens

The lens form chosen for the 34 mm lens is a Petzval type lens with only four elements, and is shown in Figure 41. This lens is similar to the lens we used previously,[128,129] except that the doublet has been separated to gain a degree of design freedom. This lens form lends itself to provide the required external stop, the $f \sin(\theta)$ mapping, telecentricity, and diffraction limited image quality over the 6.4 deg. half field at $f/2$. The lens specifications are given in Table 13 and its performance is illustrated in Figure 42. In the first optimization runs it became evident that this lens form performs much better with high index glasses such as Schott SFL6. In subsequent optimization runs, we found that the performance could be maintained by using Schott BK7 for the thick second element. This finding led to a decrease in the fabrication cost. The long aperture stop distance was very helpful because it allowed a pellicle beamsplitter to be readily inserted for alignment of the spots onto the detectors. Figure 43 shows the assembled 34mm lens in its barrel.




Figure 44. Schematic of the 106.25mm lens. The long aperture stop and image distances were particularly useful when aligning this lens in the system. Although the aperture stop and image distances are long, this system is fairly compact because the distance between the stop and the image is only 70% longer than the focal length.

106.25 mm lens

A Petzval type lens could be used for this longer focal length lens; however, the overall length of this lens type tends to be long compared to its focal length. For example, the 34mm lens is 3.3 times longer than its focal length. Since overall system size is an important consideration, we used a more compact lens form. The double-Gauss lens shown in Figure 44 permitted us to meet the all the system requirements - an external stop, $f \sin(\theta)$ mapping, telecentric image plane, diffraction-limited image quality at $f/4.2$ - all in a package only 1.7 times as long as the focal length. The specifications for this lens are given in Table 14, and the performance is illustrated in Figure 45. Initially, we designed this lens with high-index Schott SFL6 glass, which led to good performance. However, to decrease manufacturing cost, we explored the use of Schott BK7 glass and found that performance of the lens could be maintained with this lower index glass. Figure 46 shows the assembled 106.25mm lens in its barrel.

Table 14. System prescription for the 106.25mm lens. The fabrication of this lens was considerably simplified by the exclusive use of BK7 glass, one of the least expensive and easiest to work optical glasses.

	radius	thickness	glass
stop	INFINITY	33.500	
2	39.023	14.731	BK7
3	80.250	2.000	
4	31.692	14.994	BK7
5	19.492	19.017	
6	-20.451	14.507	BK7
7	-33.000	3.000	
8	100.930	11.000	BK7
9	-56.896	66.255	
image	INFINITY	0.000	




Figure 45. Wavefront aberrations for the 106.25mm lens at 852nm, for field angles of 0, 4, and 6.5 deg.




Figure 46. The assembled 106.25mm lens.

34 and 106.25 mm lenses

The 34 mm and the 106.25 mm lens work together as shown in Figure 47. The input and read beam paths from the fiber array to the detector array go through the part of the aperture stop that contains the diffraction grating, and the output path from the detector array to the fiber array go through the part of the aperture stop that contains the weak lens. The optical axis of the 106.25mm lens is offset from the axis of the 34mm lens to provide separate paths for the input and output beams. The element at the common stop in the return path is a weak lens to direct the output beams to the output optical fibers. This lens is necessary because the output fibers do not have microlenses. The distance between the object and image is 310 mm.




Figure 47. Schematic of the optical system, showing ray paths for an input beam, a read beam, and an output beam. Notice that the optical axis of the 106.25mm lens is offset from the axis of the 34mm lens to provide separate paths for the input and output beams.

Alignment features

Several techniques are used to reduce the system sensitivity to environmentally-induced misalignments. All of these techniques use careful optical and mechanical design to ensure that large mechanical changes cause only small changes in optical alignment.

First, the magnification is adjusted using an afocal pair of lenses attached to the 106mm lens. The afocal pair consists of a plano-concave, 500mm focal length lens and a plano-convex, -500mm focal length lens with their curved sides facing each other. Changing the spacing between the two lenses causes a slight change in the system magnification without appreciably affecting the quality of the wavefront. Because of the long focal length of the lenses, a large change in separation of the afocal pair causes a small change in the system magnification; a 1mm change in the spacing of the afocal pair causes a 0.2% change in the system magnification, which corresponds to a 7 μ m movement of the spots on the device array.

The location of the spots on the OE-VLSI is adjusted using a pair of 15 arc-minute wedges. Each of the wedges can be rotated independently to bring the spots onto the proper win-

dows. Because the wedges have such a small wedge angle, a large rotation of the wedges causes only a small movement of the spots; a 180 degree rotation of one wedge moves the spots by 74 μm on the chip.

One of the most sensitive adjustments of the system is the rotation of the grating. A rotational misalignment of the grating by 2.54 arc-minutes causes a spot position error of 3 μm at the edge of the chip. Because this 3 μm of spot position error is 30% of the size of the detector windows, it would result in an unacceptable degradation in system performance. To ensure that the grating rotation was as accurate as possible, the location of high grating orders were used to set the position of the grating. To take advantage of this precise rotational alignment and to simplify mounting, the weak lens was fabricated as a diffractive lens on the same substrate as the grating.

Alignment of the spots on the detectors was aided by the large separation between the 34mm lens and the aperture stop. A pellicle beamsplitter was inserted into this space to provide access for a viewport. A pellicle was used because its small thickness ensured that it would not affect image quality or spot position.

Tolerancing & Assembly

The tolerances on radii of curvature, lens thickness, and surface tilt were calculated following a method similar to the one outlined by Smith.[130] In this method, constructional parameters such as lens curvature and thickness are changed individually and the resulting image errors are tabulated. The expected image degradation of the assembled lens can be approximated by the square root of the sum of the squares of the individual image errors. Both peak wavefront deformation from nominal performance and deviation from $f \sin\theta$ mapping were used as image errors. This calculation is performed iteratively until both manufacturing tolerances and predicted image quality are acceptable.

The manufacturing tolerances we specified are the following:

- 1) Index of refraction: ± 0.0003 (H4 glass),
- 2) Radii of curvature: ± 0.02 mm,
- 3) Lens thickness: ± 0.025 mm,
- 4) Surface tilt: ± 30 arc-seconds,
- 5) Diameters: $+0.0$, -0.013 mm,
- 6) The individual surface quality should be better than $1/4$ wave; fringes must be smooth.
- 7) The individual lens wedge should be less than 1 arc-minute. The eccentric distortion introduced by lens wedges should not change the lens distortion by more than 0.001 mm over the full field of view.
- 8) Light transmission should be greater than 96% at 810 and 850 nm per objective.
- 9) The overall wavefront quality of the complete objectives should not be degraded by more than $1/10$ wave (850 nm) on-axis and four field positions separated 90 degrees at the edge of the 13 degree field.
- 10) The mechanical and optical axis should coincide within 1 arc-minute.

The index tolerance is tighter than the standard tolerance listed by the manufacturer, and represents the second-highest level of index tolerance available.[131] However, this tight index tolerance allowed the loosening of some of the other fabrication tolerances; the tilt and wedge tolerances are within what is considered standard commercially-produced tolerances. However, other tolerances, such as those on thickness, diameter, and especially individual surface quality,

required an optical shop with particularly skilled technicians[132]; these tolerances are described as “extra-precise” or “select.” [133,134,135,136] Most importantly, the 1/10 wave requirement on overall wavefront quality and distortion required special techniques such as matching specific elements to each other within an assembly, adjusting air spaces, and rotating elements about their axis; this type of challenging assembly requires an optical shop with skilled engineers.

Four 34 mm lenses and four 106.25 mm lenses were delivered from the optics shop. Testing at the optics shop included testing in a phase-shifting interferometer at 633nm to ensure no asymmetric aberrations were present on-axis. We tested the lenses’ wavefront quality in a Twyman-Green interferometer at 852nm; and we tested the boresight error by measuring the beam-wander of a 633nm He-Ne beam while rotating the lenses in a V-block. All of the lenses were tested to be under 1/4 wave of aberration in double-pass and well under 2 mrad boresight error. This excellent performance was achieved for several reasons. First, highly skilled opticians were able to meet the fabrication tolerances outlined above. Second, the as-manufactured thickness and curvature of each of the elements was measured precisely, and the lens spacers adjusted to optimize image quality. Finally, self-centering lens mounts were used to make lens assembly easier.

The first step in assembly of the optical system was to collimate the light from the input fibers. This task was performed using autocollimation. A mirror was placed in the aperture stop plane to focus the light back to the object plane. A pellicle beamsplitter was placed between the fiber bundle and the 106.25mm lens, allowing access for a CCD camera. The CCD camera was focussed on the light from the fibers, then the 106.25mm lens was moved until the light reflected from the mirror in the aperture stop was in focus in the same plane, thus ensuring that the light in the aperture stop was collimated.

The next step was to focus the 34mm lens. This focussing was performed by placing a pellicle beamsplitter between the 34mm lens and the aperture stop, allowing access for the CCD camera. The camera was first focussed on the object plane, then the orientation of the pellicle was reversed to allow the CCD camera access to the image plane. Illumination was provided by an LED through a second beamsplitter, as shown in Figure 48. The 34mm lens was then moved to bring the image of the OE-VLSI into focus at the CCD, thus ensuring that the fiber bundle plane was conjugate to the OE-VLSI plane




Figure 48. Schematic showing both the optical system and the viewport used for alignment.

The third step in assembly of the optical system was to adjust the magnification. A second CCD camera was added to the viewfinder to allow high-magnification viewing of both the top and bottom of the image plane. The spacing of the afocal pair was adjusted in increments as small as 0.003" until the spots at the top and the bottom of the field were located on the center of their respective windows. This 0.003" increment implies the maximum spot position error due to magnification error is about 0.4 μm .

The fourth step in assembly of the optical system was to adjust the rotation of the grating. Although the 1x16 grating nominally only covered 1/4 of the OE-VLSI, high-order spots were visible enough to help in alignment. The two cameras of the viewfinder were oriented so that both sides of the chip could be seen with sufficiently high magnification. The grating was rotated until the spots at both sides of the chip were aligned on the windows in the vertical direction as well as they were aligned from the magnification adjustment, about 0.4 μm . However, the alignment in the horizontal direction was unacceptable. This misalignment could have been due to inaccuracies in the lens focal length, the laser wavelength, or the grating spacing. Constructing a new grating with the grating spacing changed by 0.008% fixed the problem. The new grating placed the spots on the windows to a greater precision than we were able to measure.

The final step in assembly of the optical system was to adjust the rotation of the wedges in the pupil plane to position the spots on the proper location on the detectors. Because of the ion-implantation process used to isolate the detectors and modulators, both the sensitivity of the detectors and the contrast ratio and reflectivity of the modulators varies across the 10 μm active area of the window. Generally, the detector and modulator performance was better near the center

of the $10\text{ }\mu\text{m}$ active area and worse towards the edge. To account for this variation, the position of the spots on the windows was adjusted while monitoring the level of the output, the contrast ratio of the output, and the photojournalist from the chip for a fixed input signal on a single channel. When all of these values were good, we fixed the position of the spots on the windows by fixing the rotation of the wedges.




Figure 49. The assembled system. The figure shows, from left to right: the fiber bundle assembly mount, the 106.25mm lens with the afocal pair mounted on the end, the wedges and diffraction grating at the aperture stop, the 34mm lens, and the mount for the OE-VLSI chip.

The final assembled system is shown in Figure 49. The system fit readily into a standard electronics frame, and was mounted at a height convenient for adjustment.

Conclusions

We have described a simplified optical design for a photonic switching system. The major simplifications resulted from using a single stage design, a single fiber bundle, and pupil division beam combination. We were able to integrate alignment mechanisms and viewports to adjust the system magnification, focus, and lateral position of the image.

System considerations such as fiber spacing, detector size, and array size force constraints on the paraxial optical design of this type of system. Further constraints are imposed by the requirement that multiple output paths must be provided for each input. These constraints can be used to generate a paraxial design of a photonic switching system.

Moving from a paraxial design to a real optical design requires designing lenses. One of the most important aspects of lens design is choosing a good design form for a starting point. We discuss our reasoning for choosing our design forms, then outline the design process. In particular, we present the design of a Petzval lens and a double-Gauss lens, both of which are diffraction-limited over a 6.4 deg. field of view.

Careful system design and assembly were instrumental in the successful demonstration of a OE-VLSI based switching system. The system demonstrated reliable switching of 14 input

channels at per-channel data rate of 208 Mbit/s. Issues such as system architecture and device technology also played important roles in successful system demonstration.

Cody Kreischer of Kreischer Optics, Inc. provided technical advice and service in the fabrication of the lenses.

4.3.6 Diffractive Optics

Diffractive optic components served three functions in this free-space photonic switching demonstrator. One set of diffractive microlenses were positioned at the fiber bundle and were used to match the numerical aperture of the input fibers to the numerical aperture required at the object plane. The two remaining diffractive elements, a beam array generator and a weak focusing lens, were fabricated on the same optical substrate and positioned at the system pupil. The beam array generator multiplexed each incoming beam into a linear 1x16 spot array. The weak focusing lens was used to assist in coupling the modulated output beams into the multimode fibers. The position of the two pieces is shown in figure 47.

All diffractive optical components were designed to operate at 852nm wavelength. All components were fabricated on fused silica substrates that had an antireflection coating on the side opposite the diffractive features. The round substrates were 25mm in diameter and approximately 1mm thick.

The diffractive microlenses were designed to have a focal length of 1.0515mm and a window diameter of 250 μ m on spacings of 250 μ m. A linear array of about 80 lenses were designed although only a fraction were used by the existing fiber array. Three etch steps were used to create an eight phase level element on the fused silica substrate which yields a theoretical diffraction efficiency of about 95%. The minimum feature size was limited to 1.0 micron.

The weak focusing lens was designed to have a focal length of 5061.3 mm. The center of the lens was offset from the center of the round substrate by 8.93 mm. Three etch steps were used to create the eight phase level element. The lens design was windowed by a rectangular area of size 10mm by 9mm that was located on one half of the substrate's surface.

The beam array generator was designed to create a linear array of 16 uniform intensity spots when operated at 852nm using the 34mm focal length objective. This required a 723.6 μ m period. (Note: the grating design has suppressed even-numbered orders, leaving only odd-numbered orders to create the 16 spots.) The design's diffraction efficiency is 90.7% with a theoretical intensity deviation of under 1%. The 1x16 design was created from a basic 4096 phase cell pattern although the minimum contiguous feature is 14 cells corresponding to a minimum feature size of 2.5 microns. The design was created using a cell-based algorithm developed by a member of the Bell Labs team.[137] The array generator fit within a 10mm x 9mm area on the opposite half of the substrate with the weak focusing lens. Although the relative spot intensities and efficiency were not measured with these sets of elements, it is expected (based on our previous work in characterizing diffractive elements) that the efficiency closely matches the design value reduced by about 4% reflection loss from the diffractive surface and the intensities vary by no more than +/-3% for the 1x16 array and vary by about +/-7% for the larger 1x68 array described next.

An additional beam array generator was also designed and fabricated in order to illuminate the full array and test all modulators. In order to illuminate four sets of modulators which include

gaps between modulators set, a 1x68 linear array generator was designed. The two phase level pattern had a theoretical diffractive efficiency of 79.7% and a intensity deviation of about 2%. The less efficient two level design was used since it was easier to fabricate. The previously described weak focusing lens was also fabricated on the same substrate.

4.3.7 Mechanical Design

To be competitive with electronic switching technologies, photonic switching systems must have stability that rivals that of electronic systems, requiring little or no intervention over time scales of years. This paper describes progress towards the goal of stable optical interconnects. We have built an extraordinarily stable free-space optical interconnect mounted in a standard electronics frame; the system operated successfully over a wide temperature range for three days and required no realignment after shipping. Robust optomechanical design played an important role in the successful operation of the system. The role of kinematic mounting principles, self-centering lens mounts, materials selection, and long-lever arm adjustments are described. Vigorous shaking of the system did not affect its bit-error rate - measured to be less than $10E-12$ on a single channel.

Introduction

Current, commercially-available electronic switching systems present high benchmarks for optical switching systems to meet. At the top of the line of ATM switches, for example, is AT&T's Globeview 2000 switch.

The Globeview 2000 is an all-electronic ATM switch which offers aggregate throughput of 20 Gbps. It can operate in a typical business-office environment, which has a temperature variation from 18 - 24 deg. C (65 deg. F - 75 deg. F), and a humidity variation of 5% - 60%. It requires no special vibration isolation, and some systems have survived earthquakes. It is shipped to the customer premises in standard trucks, and requires only a few hours to become fully operational. It is designed to operate for several years with a cell-loss probability of less than 10^{-9} with only occasional maintenance.

A direct comparison between electronic switching systems and optical switching systems is not really fair. For example, optical switching systems can use many electronic reliability components developed for electronic switching systems. Nevertheless, some issues such as thermal stability, shipability, and set-up time must be addressed if optical switching systems are to replace electronic ones. This paper describes some of the efforts we have made in improving the ruggedness of optical switching systems.

The system described here, System 6, is a rack-mounted system that successfully switched two channels of video encoded into 155 Mbps ATM streams for three days at a conference showroom at the National Communications Forum.[14] The temperature in the room varied from around 13 degrees C (55 deg F) to around 26 degrees C (79 deg F). The system required no re-alignment after being shipped in a rental truck.

This extraordinary stability is due to a high-quality DBR lasers, robust design of the hybrid FET-SEED modulator[73], tolerant system architecture, and rugged electronic design of the laser modulators and receivers, as well as robust optomechanical design. This paper describes only the optomechanics of our system and some of the guiding optomechanical design principles.

SYSTEM ARCHITECTURE

Previous systems built by our group have been multi-stage systems, with the image of one modulator array being cascaded onto a second modulator array[104,117,118]. Some of these systems have had as many as six stages. Advances in system architecture allowed us to design this system with only one stage.

Reducing the number of stages was the most important aspect of system design to improve system durability. The description of the architecture is given in section 4.3.2.

FRAME MOUNTING

The system is entirely contained within a standard electronics frame, which consists of 12 beams formed into a box (80x80x300cm), which is mounted on casters for easy transportation. We chose this structure purely because it is a standard method for mounting switching systems such as AT&T's Globeview 2000 electronic switching system.

These standard electronics frames are generally sold as modular systems. Various sizes are available, and various options are available for each size. For example, our frame was manufactured by Schroff GmbH. We chose our frame from a variety of standard frames with a variety of standard shelving units, doors, and panels.

The mechanical properties of these frames in terms of rigidity, vibration isolation, and thermal stability are far from the comparable properties for optical tables. Since all of our previous systems were designed to operate on optical tables, we had to re-examine many of our optomechanical design schemes so that the required optical alignment tolerances would be maintained.

The electronics for our system required six standard shelves within the rack. These shelves added some rigidity to the structure, and additional stability was provided by the standard doors and sides which we placed on the rack. However, even with these additional measures, the cabinet shakes several centimeters with only a slight shove. The assembled frame seems to be less sensitive to high frequency vibrations; no ringing is evident when the frame is tapped with a hammer.

The optical system is placed in a steel sheet-metal box, which is bolted between two of the electronics shelves. The bolts are placed through standard 0.75" rubber vibration mounts which provide 81% isolation at 1kHz, providing some high-frequency vibration isolation for the optical system. The box was placed at eye level (about 4 1/2 feet from the ground) for ease of alignment and for display purposes. Placing the box at this level requires extra care to be taken to assure that the stray light escaping from the system is within safety standards.

All of the optical system is contained within a second, plastic box which slides into the frame-mounted metal box. Access holes are provided for the fiber bundle, electrical control to the chip, and the alignment viewport.

MECHANICAL MOUNTING

The lens, SEED, fiber bundle, and other optics are mounted in a beam structure, which is mounted inside the second, plastic box, and is supported near the Airy points of the beam. The Airy points of a beam are defined as the points at which the endpoints of a loaded beam have equal and parallel deflections, as shown in Figure 50[138]. The beam structure measures about 60x70x370mm and is mounted in kinematic mounts of a ball-and-cone type.

Figure 51 shows an exploded view of the optomechanical mounting for the system.

The entire beam structure is made of 6061 Aluminum alloy. Aluminum is a good material to use because it is one of the easiest materials to machine. However, there are sound optomechanical design reasons to choose aluminum as well. Optomechanical mounts such as this one must be both lightweight and stiff. A good measure of stiffness is the amount of self-weight deflection, which is proportional to ρ/E (density/Young's modulus) Table 15 lists the density and elastic modulus of several materials used for optomechanical systems. The table shows that Aluminum has a stiffness-to-weight ratio that is as good as any likely material to be used. Beryllium, for example, has a superior stiffness-to-weight ratio, but has serious drawbacks, including very high cost and poisonous fumes.




Figure 50. Illustration of the concept of Airy points, showing a self-weight-loaded beam supported a) inside the Airy points, b) at the Airy points, and c) outside the Airy points. For the beam mounted at the Airy points, the deflections of the end points are equal and parallel. For the other two beams, the deflections of the end points are equal but not parallel.

6061 aluminum alloy is a good material to ensure long-term stability of the structure because it is heat-treated to reduce microcreep. Also, it has good thermal conductivity to reduce misalignments from thermal gradients. Finally, because it has such excellent properties, it is used in many precision applications, so its properties are well-controlled and it is widely available at a reasonable price.

Table 15. Important opto-mechanical properties of several mounting materials.

	Density (ρ) $\text{kg/m}^3 \cdot 10^6$	Elastic Modulus (E) GPa	ρ/E ($\text{m}^{-1} \cdot 10^9$)
Aluminum 6061-T6	2.71	68.9	386




Figure 51. Optomechanical schematic

	Density (ρ) $\text{kg/m}^3 \cdot 10^6$	Elastic Modulus (E) GPa	ρ/E ($\text{m}^{-1} \cdot 10^9$)
Steel 1040	7.84	200	385
Copper OFHC	8.89	117	1745
Beryllium I-220	1.85	303	59.9
Invar 36	8.0	145	541
Graphite Epoxy	1.8	293	60.3
Silicon Carbide SXA24-T6	2.91	117	60.3

The beam-and-pedestal structure chosen for mounting the lenses provides an excellent stiffness-to-weight ratio, while still providing access to the system for alignment. Stiffness-to-weight ratio is important for reducing misalignments from self-weight deflection. A useful model for examining self-weight deflection is a self-loaded, edge-supported beam; for this condition, the total deflection is inversely proportional to the moment of inertia[139]. Because of its larger moment of inertia, a beam-and-pedestal structure has about 28% less self-weight deflection than a comparable tube structure.

The pedestals are mounted to the baseplates using 8/32" bolts, separated by 1/2", giving 28 bolts per pedestal. This close spacing of the bolts allows a complete coupling between the pedestal and

the plates, allowing the entire assembly to act as a single beam. The optical components are held in the pedestals using a clamp structure that was cut into the pedestal using Electric Discharge Machining (EDM). This special type of machining was necessary to make the relatively delicate clamp structures in the same piece as the relatively bulky pedestals. Also, because EDM machining has no tool wear, it allows the bores to be made to a very accurate inner diameter. All of the pedestals were sized and bored simultaneously to ensure accurate alignment of the final system.

The fiber bundle is held in the pedestal using a clamp structure with reference pins for accurate registration. To ensure roundness, the two pieces of the clamp were machined simultaneously instead of being machined as one piece then cut in half. A pin on the clamp holds fixes the axial position of the fiber bundle. The fibers are held in a lightweight plastic strain relief that isolates each fiber in a comb structure. The entire assembly can be rotated around the optical axis for rotational alignment with the SEED.

The construction of the fiber bundle is described in another paper[140]. This technique places the location of the fiber cores with an accuracy of about 1 micron - better than the core-cladding accuracy. Silicone in the fiber bundle holder provides a first measure of strain relief. A second level of strain relief is provided by a plastic housing which clamps to both the aluminum fiber bundle holder and each of the fibers. The fibers are jacketed in standard plastic jackets and terminated with either ST or FC connectors.

A diffractive microlens array is placed on the end of the fiber bundle. Special alignment marks on the microlens array and on the fiber bundle mount allow accurate measurement of the microlens alignment. The microlenses are positioned using manual micropositioners. Once aligned, the microlens array is temporarily held in place using a single drop of UV-curing cement, placed at the edge of the microlens array. Finally, the microlenses are held in place by a clamp which holds the edge of the microlens array to the fiber bundle holder.

Both the collimating lens and the focusing lens are in precision barrels. The lens elements are held with self-centering mounts, as shown in Figure 52[141]. The convex surfaces are held using tangential mounts and the concave surfaces are held using sharp-corner mounts. The lens barrels can be moved by hand axially to focus and can be rotated about the optical axis to cancel some resid-

ual axial coma.




Figure 52. Illustration of a self-centering lens mount. As axial force is applied to the lenses by a retaining ring, the lenses tilt until their precision optical surfaces are in ring contact with the ends of the spacers, which are machined to an angle tangent to the lenses. This mounting scheme removes the need for precision edges to be ground on the lenses.

.....

ALIGNMENT AND ADJUSTMENTS

Several techniques are used to reduce the system sensitivity to environmentally-induced misalignments. All of these techniques use careful optical and mechanical design to ensure that large mechanical changes cause only small changes in optical alignment.

First, the magnification is adjusted using an afocal pair of lenses attached to the 106mm lens. The afocal pair consists of a plano-concave, 500mm focal length lens and a plano-convex, -500mm focal length lens with their curved sides facing each other. Changing the spacing between the two lenses causes a slight change in the system magnification without losing the quality of the wavefront. Because of the long focal length of the lenses, a large change in separation of the afocal pair causes a small change in the system magnification; a 1mm change in the spacing of the afocal pair causes a 0.2% change in the system magnification. Each of the lenses in the afocal pair is mounted in a barrel around its edge in a split nylon ring that is tapered so that the lens is held in place by a threaded retaining ring. One of the afocal pair barrels is clamped onto the end of the 106mm lens barrel. The second afocal pair barrel then screws onto the first afocal pair barrel; the separation between the two lenses is adjusted by unscrewing the second lens, and is fixed by fixing a threaded retaining ring in the proper location.

The location of the spots on the SEED is adjusted using a pair of 15 arc-minute wedges. Each of the wedges can be rotated independently to bring the spots onto the proper windows. Because the wedges have such a small wedge angle, a large rotation of the wedges causes only a small movement of the spots; a 180 degree rotation of one wedge will move the spots by 74 microns on the chip. The wedges are cemented into rings that have holes and notches for easy rotation. The position of the wedges is fixed by the same type of clamping mechanism that fixes the position of the

lenses.

One of the most sensitive adjustments of the system is the rotation of the grating[142]. A rotational misalignment of the grating by 0.001 degree causes a spot position error of 3 microns (30% of the window size) at the edge of the chip. Careful optomechanical design was necessary to achieve this accuracy. The grating is cemented in an aluminum mount using a thin layer of UV-curing cement around the edge of the glass; long-term stability of the grating rotation is aided by ensuring even illumination of the cement during curing. The grating mount has two pins on it for adjusting rotation. Two screws in the pedestal push against these pins to adjust the rotation of the grating. Precise rotation of the grating is possible because a large rotation of the adjustment screw causes a small rotation of the grating. The rotation of the grating can be adjusted to a high degree of accuracy by examining the position of the higher grating orders. The position of the grating is fixed using a clamp.

The SEED mount is an aluminum cylinder that mounts into a pedestal. The cylinder has a lip to fix the axial position of the chip. The chip is mounted on a mesa on the face of the cylinder using silver paint; the mesa allows the flex circuit, which is used for electronic control of the chip, to have relatively large terminating resistors placed close to the chip. The flex circuit, which brings power and electronic controls to the chip is cemented to the mesa, then clamped to the edge of the pedestal. Finally, the flex circuit is wire-bonded to the chip.

4.3.8 Fiber Bundles

Motivation

The optoelectronic device technology for photonic switching is at a very advanced state and appear very promising for including into products. However, interfacing to these advanced devices requires a way to get the optical channels into the system. Fiber bundles would be an excellent way to bring this data into the system.

Our ideal fiber bundle would be:

- 64x64 fibers
- fiber cores positioned within about 1 micron
- communication-grade fibers, single mode at 852nm
- 5.12mm x 5.12mm or smaller

Challenges:

- Fiber Handling

Handling 4,000 fibers will require special methods. Dealing with the fibers in ribbons seems to help in fiber handling, but nonuniformities within the ribbons seem to be a problem.

- Metrology

Measuring the position of all 4,000 fibers will require an automated system. Fitting those 4,000 points to the best-fit grid, and determining patterns in the deviation from that grid will require some non-trivial programming.

- Terminations

The face of the fiber bundle will have to be polished flat. Suitable epoxies for holding the fibers and polishing methods will have to be determined.

All of the 4,000 fibers will have to be terminated in some way. Terminating a fiber with an SC-type connector will cost at least \$20/fiber and take a space of 12x20mm, for a total cost of \$80,000 and a total area of 768mm x 1280 mm. Clearly, we would like to reduce these values. Twelve-fiber multi-mode ribbon connectors are available in a 20x14mm package; terminating an entire fiber bundle with a similar (as-yet unavailable) single-mode connector would require a space of at least 380mm x 266mm.

- Automated assembly

Commercial viability is often mentioned as a motivation for automated assembly. However, for fiber bundles of this size, some form of automated assembly will be required even for laboratory demonstrations. For example: Stripping and cleaving a single fiber requires at least two minutes, so stripping and cleaving 4,000 fibers with current, in-lab processes, will require four weeks of five, eight-hour days (allowing for a 10-min break every hour).

Fabrication of Fiber Arrays

In this section, we describe a technique for assembling fiber arrays as needed in optical computing and photonic switching. A 4 x 8 array was manufactured with fiber ends to within 1.5 μm from their ideal position and to a pointing precision of 30 arc-minutes.

Optical fiber bundle arrays are simple to conceptualize but their fabrication has proven to be difficult when high precision positioning is required. Several fabrication techniques have been reported in the literature; for example, in the approach presented by Miller [143] a 2-D array of fibers was made by stacking a number of linear arrays of fibers supported by grooved spacers. These spacers were manufactured by the precise etching of both sides of a silicon wafer, and a polishing operation was involved after potting all the fibers in place. Recently, this technique has been used by Danzer, Kipfer, Zurl, Lindolf, and Schwider [144] to assemble a fiber array with a maximum positioning error of 10 μm . In another effort an alignment-free assembly technique has been developed by Sasaki, Baba, and Iga [145] where fiber end positioning was accomplished to within + 8 μm . In this technique an array of fiber sockets with centering plugs were micro-fabricated to achieve fiber self-centering insertion and to expedite assembly. Koepf and Markey [146] have reported a technique involving arrays of precision holes in substrates to insert and locate optical fibers with a standard deviation of 12.6 μm . Arrays of precision holes to position fibers have also been used by Basavanahally [147] and by Proudley, Stace, and White [148] to manufacture fiber bundles. The insertion of fibers in microferrules and their stacking to create a 2-D fiber array has even been realized [149] with a mean fiber positioning error of 3 μm . A common feature of these techniques is that fiber positioning is accomplished by referencing the fiber to a mechanical jig and this limits the ultimate precision attainable. We considered these assembling techniques but were not satisfied with their inherent errors or the processes involved. Thus we developed the assembly technique that we present which suits our photonic switching development.

The basic idea around which our technique revolves is the high precision and individual positioning of each optical fiber on a substrate of large and low precision holes. Instead of relying on a mechanical substrate to reference the fiber we obtain positioning accuracy by individually

locating each fiber core in the correct position. The positioning accuracy is achieved by referencing light coming from the fiber core to a lithographically made array of annuluses or doughnuts on a transparent glass substrate. Once a fiber is positioned, it is bonded in place by curing W cement around it. The array of holes serves as the foundation for the fiber array, to achieve uniformity in fiber pointing, and to stress relieve the optical fibers. To assemble a fiber bundle a hole array substrate is filled with UV curing cement and an array of centering doughnuts is registered on the top surface of the hole array. The centering doughnuts face the array of holes and are centered with respect to the holes. The assembly formed by the hole array and centering doughnuts substrate is placed on a microscope. Prior to the fiber insertion and alignment steps the fiber coating is stripped off and the fibers cleaved. With the aid of a fiber manipulator a fiber is grabbed and inserted through the corresponding hole in the hole array. The light that comes from the fiber core and some TV cameras make this insertion easy. The fiber can pivot on the lower rim of

the hole so that it can be aligned by horizontally moving the fiber manipulator. Then the fiber is brought into contact with the doughnut substrate. For this delicate step a piezo-electric driven mechanism is used. Once the alignment has been done a pinhole is positioned on top of the fiber core to only allow the cement in the hole of interest to be exposed to W light. At this time the fiber position is verified and then a W light gun is shone on top of the pinhole to cure the cement and bond the fiber to the hole and substrate. This bond holds the fiber in position and also helps to stress relieve the glued fiber end from the rest of the fiber. Other fibers are glued in a similar fashion and in an orderly way to avoid interference between the optical fibers as they are handled and moved under the hole array substrate. Physical interference between fibers is avoided by slightly bending the fibers after gluing so that all the fibers are moved out of the way of the next fiber. The elapsed time between each fiber alignment and gluing is approximately 10 minutes. Using the technique described we manufactured several fiber arrays. A recent one has 32 single mode fibers at 1300 nm in a 4 x 8 array and a distance of 500 μ m between fiber centers. The fiber ends were located to within 1.5 μ m from their ideal position and to a pointing precision of 30 arc-minutes. We have developed a technique for assembling fiber bundle arrays with improved fiber positioning accuracy. The technique does not depend on fiber core concentricity, on fiber dimensional uniformity, or on a final polishing step. Except for the use of lithography for fabricating the centering doughnuts all the equipment necessary to implement this assembly technique is simple and is easily obtained. We use active fiber alignment to center illuminated fiber cores within lithographically made referencing doughnuts to achieve micrometer precision in fiber end positioning. In addition, the array of centering doughnuts makes the detection of alignment errors straightforward and limited by the microscope resolution. The assembly method is very well suited for the manufacture of small arrays and for the development of optical computing and photonic switching systems where only a few fiber arrays are required but with changing requirements. Some drawbacks of the technique are the assembly time required to construct large arrays, and that it takes only one broken or mispositioned fiber to ruin an entire array. This latter drawback is common to all techniques. Thus fiber assembly reliability and yield are important figures of merit to compare assembly techniques. For future work we plan to automate the technique to decrease assembly time and increase reliability. We also plan to manufacture larger fiber arrays, study fiber array repairability, dimensional stability, and incorporate arrays of microlenses to match the numerical aperture of the fiber bundles to that of our photonic systems.

System 6 Fiber Array

To fully access the switching chip, the fiber bundle requires 256 inputs data fibers, 16 "read" fibers with continuous wave inputs to read the state of the modulators, and 256 data output fibers. Because this is a time consuming task (but necessary for real systems), we chose to build a fiber bundle to operate a portion of the switch. The fiber bundle consists of 16 input fibers, 4 read fibers, and 16 output fibers as shown in Fig. 53. The input and read fibers were 850 nm single mode fibers and the output fibers were multimode fibers. The fibers are designed to provide input beams to inputs 0, 4, 11, and 15 from each of 4 switches in a column. Including the read beam fibers, the single mode fiber spacing imaged within a switch is 1 mm. That is, every fourth input has a fiber. However, the spacing between the last input from one of the switches and the first input to another is only 250 μm . The four outputs per 16 x 16 switch are designed to collect light from outputs 2, 6, 10 and 14, also for each of four switches in a column. Thus the output fibers were also on 1 mm centers. The geometric center of these four fibers was offset by 125 μm from the modulator read fiber to account for the fact that the 16 orders of the grating when imaged back onto the fiber bundle are on odd multiples of 125 μm from the zero order.


4.3.9 Input Lasers

Because all the input lasers and read lasers pass through the same binary phase grating, all lasers must have their wavelengths approximately equal and stabilized. We chose distributed Bragg reflector (DBR) lasers at 852 \pm 1 nm for these lasers. Six of the 18 lasers had thermo-electric coolers to stabilize their wavelengths. The lasers were connected (pigtailed) to 850 nm single mode fiber, which was then connected to the fibers of the bundle using a fiber to fiber female connector. The input lasers were electrically driven by a commercial laser driver chip that accepted emitter coupled logic (ECL) level inputs and provided adjustments for offset (threshold) and maximum output power.

Extensive characterization of individual lasers was performed. The lasing wavelengths were between 850.8 and 853 nm. The shift in wavelength when modulated at 200 Mb/s was less than the 0.2 nm resolution of the spectrometer. Wavelength shifts could not be introduced when back reflections were introduced.

One of the most crucial issues is the temporal waveform under large signal modulation. We required a high contrast ratio from the lasers to ensure an adequate noise margin on the single ended receivers on the switching chip. A reasonable specification was 13 dB (a factor of 20) which means the ratio of the optical power of a high signal to the threshold and the ratio of the threshold to the optical power of a low signal are both slightly greater than four. This must include any variations in output power over time, errors in positioning and spot size that cause the coupling into the detector to vary both spatially across the array and over time, and any changes in threshold of the receivers themselves, again both spatially and in time. Changes in optical power, both in the high and low state, cause delay variations or pulse width distortion in the receiver. Thus, it is not enough to have the optical powers above or below the threshold (depending on the input state), they must exceed the threshold by a significant amount so as to cause only minimal pulse width distortion.

Because of the high contrast ratio required, the lasers were biased below threshold. Biasing below threshold, can lead to excessive turn-on delay, which also leads to pulse width distortion in the



- output fibers (multimode)
- input fibers (single mode)
- read laser fibers (single mode)

Figure 53. System 6 fiber bundle showing the locations of the input, output, and read fibers

system. To keep the turn on delay below 100 ps, the lasers needed to be biased above 89% of the threshold current. However, to keep the contrast ratio above 13 dB, the lasers needed to be biased below 1.002 times the threshold current. Long term power stability is an issue when biasing this close to threshold. We found qualitatively that the stability is good enough for operation of the switch at 200 Mb/s.

4.3.10 System Control Software

The function of the application software was: to characterize the operation of the system 6 chip

before integration into the system, to configure the sixteen individual switches on the chip during system operation and testing, and, to run a demonstration program illustrating the capability of the switch to route a pair of digitized video signals.

The control application was written in C and ran on an Apple Macintosh IIcx computer. The fundamental control signals and data signals for the system were produced by a Tektronix HFS9000 stimulus generator that communicated to the Macintosh via the GPIB bus and to system 6 via individual coaxial cables.

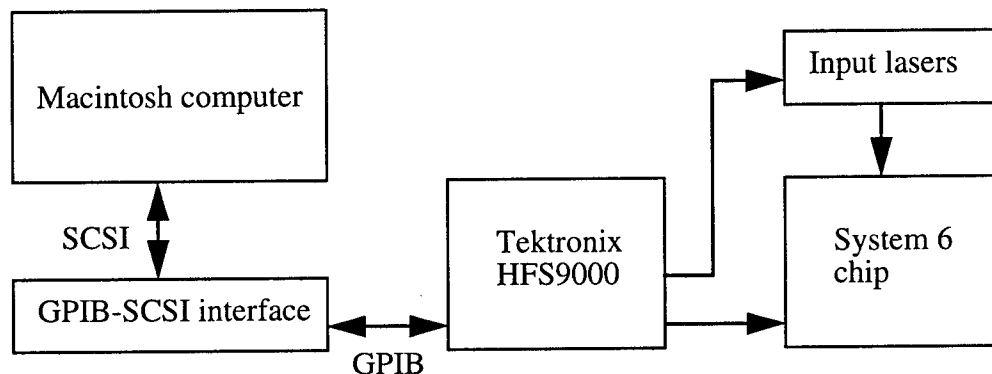


Figure 54. System 6 control diagram.

The Tektronix HFS9000 controls about 30 signals. There are 8 data lines for the input lasers, 4 idle channel lines, 16 encoded control lines for setting switch configurations, a clock line for latching the serial control information, and a load shift line for signaling the chip to switch configurations to the new control specification. A few other assorted signal lines are also available for triggering the oscilloscope or for performing electronic diagnostics. Each of the channels on the HFS9000 was configured with about 64Kbits of memory providing the ability to send streams of preset data values. Typically, the HFS9000 would cycle through a predetermined number of bits composed of control information for switch configuration followed by an arbitrary length packet of data. Due to the limited speed of the GPIB bus, new configurations parameters might require several hundred milliseconds to load, thus the HFS9000 was typically halted throughout a data transfer. Pseudo-ATM-like behavior, however, was achieved by presetting a series of 16 ATM packets in HFS9000 memory which would reconfigure the switches approximately every 2 μ s.

The application program is a menu driven program that allows the user to specify several operational parameters of the photonic system. For example, the user can separately set the data values of the input channels, set the operating channel bit rate, set the various switch configurations, run an automated video switching demonstration, and control various system diagnostics.

Figure 55 shows the display for setting the routing through the individual switches. There are several features illustrated in the figure. The dominant section filled with "radio" style buttons allows the user to connect any one of the sixteen input channels or the idle channel to an output modulator. The radio button term refers to the fact that one and only one input/idle channel may be selected. Because there are 16 separate switches on the optoelectronic chip, a 4x4 radio button

matrix on the left side of the display selects the appropriate switch. Since the all configurations are stored in memory on the computer, the channel/modulator display is immediately updated with the last configuration when a switch is selected.

Control Inputs

SWITCH CONTROL

ATM Cell ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐

☐ ATM Enable

Output Modulators

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
1	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
2	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
3	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
4	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
5	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
6	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
7	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
8	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
9	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
A	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
B	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
C	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
D	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
E	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
F	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
IDLE	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

☒ ☐ ☐ ☐

SWITCH

Cancel

OK

Figure 55. Display for setting one of the system 6 switch configurations.

The row of radio buttons along the top of the display is only used when the system is operated in an ATM-like mode, that is, when the switches are reconfigured after each 53 byte cell. When the ATM Enable is not checked, all switch configurations are held static during operation. When the ATM Enable is checked the switches are dynamically reconfigured in a sequential manner. The routing characteristics can be configured for each of sixteen individual and independent ATM cells.

The display for the free-space photonic switch demonstration is shown in figure 56. In the demonstration two digitized video signals are routed through one of the sixteen switches on the chip. The row of four switch icons at the bottom of the display show the four basic switch configurations that are demonstrated: A->A', B->B'; A->B', B->A'; A->A' and B'; and B->A' and B'. The switch icon in the center of the display shows the current switch configuration, while the AT&T and Lucent icons on the right are updated to dynamically illustrate the switching. When the Auto Scan box is checked, the demonstration sequentially chooses each of the four configurations changing at the rate of about once every four seconds. The viewer is also able to watch the video signals being exchanged on two large video monitors. Typically, one input channel is connected to a

video camera in the lab while the second input video channel is connected to a continuously running laser disc player.




Figure 56. Display for the automated switched video demonstration.

4.3.11 Experimental Results

The fiber bundle provides 16 data inputs and 4 inputs interrogating 4 of the 16 switching chips. Because two of the input paths were inaccessible due to a "stuck" control bit, fourteen lasers were connected to the input fibers. These lasers were modulated with fixed patterns from a digital word generator. Four lasers were connected to the four read fibers and these lasers were not initially modulated. The fiber bundle has 16 output fibers, and the 16 outputs from these fibers were connected to a programmable fiber switch to facilitate faster testing. Since there are 16 output fibers and each output fiber can select data from one of four input fibers, there are 64 accessi-

ble paths through the network, including the two inaccessible fibers (56 not including those).




Figure 57. Outputs from each of the 16 output fibers, each with four inputs selected at 200 Mb/s when configured as a space switch.

Three tests were performed to verify operation of the switch. First, we operated the switch as a slowly reconfigurable space switch, with a data rate of 200 Mb/s and a control rate of 25 Mb/s. Laser powers for the individual data input channels were adjusted while the data rate was turned up to 300 Mb/s. At this data rate, all paths were operational immediately following adjustment, although the pulse width distortion was poor. The paths did not remain operational overnight, because of variations in laser power over time. We think these variations were caused by the poor electrical connectors that provided a precision voltage reference used in the laser driver circuit, although temperature induced threshold changes could also have caused variations in power. At 200 Mb/s, the switch had plenty of margin, and no measurable change in pulse width was observed for several days. The data from the output fibers is shown in Fig. 57 for the 64 paths through the network for inputs of a fixed pattern. Bit error rate was measured with a 2^{23} pseudo-random sequence supplied to one of the input lasers, while the other lasers had the fixed pattern. Error rates below 10^{-13} were observed.

As we just stated above, variations in optical detected power, both in time and spatially across the array limit the bit rate of the chip in the system to a rate significantly less than the rate at which one can operate a node within the array. These variations in power cause pulse width distortion, which if bad enough causes bits to be missing. We have measured the pulse width distortion as a function of both optical power and power supply voltage for receivers near the four corners of the array [73]. A rough estimate of the result is that a 50% increase or decrease in power causes a 0.75 ns increase or 1.5 ns decrease in the pulse width at 200 Mb/s. Also, the variation across the array is nominally less than 400 ps. If we neglect this variation, we can get some indication of the optical detected power variations across the 64 accessed inputs by looking at the pulse width in Fig. 57, and compare that against the power dependence that we measured for the individual receivers on the chip. After compensating for the difference in cable lengths of the data in Fig. 57, the measured pulse widths are 4.9 ns \pm 1.1 ns. From this, the detected power across the array was 35-158 μ W. These variations in power include not only the difference in laser power, but differences in the percent of light coupled into the detector windows as well. Other factors that affect this measurement include pulse width distortion and skew of the lasers themselves (with respect to one another), variations in threshold across the array, variations in cable delay among the cables of a given length, and variations in fiber length.

In a second experiment, the digital word generator supplied 576 bit (72 byte) cells at 208 Mb/s representative of cells that could have been supplied by the input interface unit. This "pseudocell" consisted of an 8 bit preamble (00000101), an 8 bit cell number from 0 to 16 differentially encoded (for example: 01011001 for cell number 2), 69 bytes of a data consisting of a repetitive cell representing the input channel number, also differentially encoded, and an 8 bit postamble (0101000). The differentially encoded data was used because it is easier to visually identify compared with non-differential data with 8b/9b encoding. The 69 bytes is more than enough for the ATM cell (53 bytes), 8b/9b encoding (<7 bytes), and additional overhead functions. The data repeated every 16 cells.

Fig.58a shows one of the detected output from one fiber at the end of the 15th cell and the beginning of the 0th cell. The transition between the two cells is quite evident, looking much like an extra bit. This transition did not occur unless the individual multiplexer reconfigured between cells. This rules out current spikes on the power and ground leads during the transfer of data between the primary and shadow memories as the cause. Thus, the likely source of this visible transition is overshoot from the receiver output when it is switched from being disabled to being enabled.

Figure 58b shows the outputs from all 16 fibers for the first four cells. Since the fiber bundle could supply only four of the inputs, we chose to cycle the data through these four rather than through all sixteen. Thus, other than the cell number, the data is repetitive every four cells. Careful inspection showed that the patterns are correct, recalling that the inputs to fibers 15 and 16 were unavailable.

No special synchronization was performed on the input fibers other than the fixed delay added by the signal generator to the 10 foot cables to give them delay equal to the 15 foot cables. The resolution of the delay was 200 ps. The control signals were delayed relative to the data inputs to set the reconfiguration of the switch in the center of the guard band interval as illustrated in Fig 58. The delay variation within the array was $\sim \pm$ 1.5 ns and all channels reconfigured well within the 40 ns guard band.




Figure 58. Outputs from (a) one channel and (b) 16 channels for the switch operating as a time multiplexed space switch at 208 Mb/s. The outputs are shown during reconfiguration from one cell to the next. The data in (b) shows the first 4 time slots of 16.

The output interface unit in a real system was to have removed the guard band from the data, so that the "glitch" during the transition would not have caused any loss of data. Since these units were not built, it was difficult to measure bit-error rate during pseudo-ATM operation. To remove the transition pulse, the detected output from one channel was summed using a resistive

splitter/combiner with a corresponding negative pulse from the word generator synchronized to the transition pulse. Another channel of the word generator was programmed with the expected pattern and connected to the bit error rate detector which compares the reference pattern to the detected output (with the removed pulse). Bit error rates below 10^{-13} were measured for this channel in overnight testing.

To operate and diagnose a larger section of the array, the 1 x 16 phase grating was replaced with a 1 x 68 grating. A 1 x 68 grating was used rather than a 1 x 64 grating because the array had "gaps" in between the four sections of the array. This grating fanned out each of the 14 inputs across the entire chip in the horizontal direction. Thus, there were $64 \times 14 = 896$ receivers with incident light, of which 256 were selected, and all 256 optical modulators had incident read beams and were modulated with incoming data. To access these 256 outputs, we used the video sampling oscilloscope [108]. A pellicle in the output pupil imaged the light from all 256 modulators onto a high resolution camera. A 40 mm focal length triplet imaged the array onto the camera. This provided an ideal magnification of the modulator array nearly fully filling the CCD array. The video sampling oscilloscope uses a pulsed modulator read beam, synchronized with the modulator drive function (in this case the word generator that drives the data input lasers), but with a slightly different (~ 1 Hz) frequency to "strobe" or sample the repetitive optical output. The word generator that supplies the data inputs was phase locked to one signal generator set to 160.0000001MHz and the four read lasers were connected to pulse generators, which were triggered by a second signal generator set to 20.000000MHz. The factor of ~ 8 allows one to sample an eight bit word with a video bit rate of 1Hz. The pulse width of the read laser was ~ 1.5 ns which did cause some "rounding" of the waveforms compared to our early experiments that used a 200 ps FWHM pulse, but it was sufficient for determining the extent to which the array was operating. Fig. 59a and 59b show the array operating under two different control and bias levels. Receivers in alternating columns had different feedback resistors and thus different thresholds or sensitivities. At 160 Mb/s, we could not get both receivers to operate concurrently under any bias condition. This would have been possible with greater optical powers and/or at a lower bit-rate, both of which effectively increase the dynamic range of the receivers. In Fig 59a, the outputs in the odd columns selected receivers with incident light and valid data and the even columns selected receivers with no incident light. The opposite is true in Fig. 59b where the Vdd supply voltage was adjusted (from 5.22V to 4.75V) to allow operation with the less sensitive receivers.

The measured spread in wavelengths of the input lasers was from 850.8 to 853.0 nm. The design wavelength of the 1 x 68 grating was 851.7 nm. The differences in wavelength cause the spacing of the spots to change, which causes the spots near the edge of the array to be misaligned with respect to the optical windows. The best performance was obtained with longer wavelengths than the design wavelength, this discrepancy could be due to the accuracy of the wavemeter used to measure the wavelength and the degree to which the focal length of the objective lens is known. Four inputs that were fully operational across the array had wavelengths of 852.3 - 852.8 nm. The difference in wavelength of the lower wavelength inputs (850.9) from the optimum (852.5) causes a spot shift of ~ 5 μ m near the edge of the array. This is enough to reduce the power coupled into the detector by almost 50% near the edge of the array. Also, distortion of the lenses themselves, errors in the positions of the fibers, and aberrations such as field curvature potentially cause even less light to be coupled into these detectors near the edge of the array. Thus, only some of the inputs were operational across the full field of view. However, all 14 inputs were operational across the 16 nodes that the experiment was originally designed for, where the misalignment for

the same wavelength variation was a maximum of $\sim 1.25 \mu\text{m}$.




Figure 59. Outputs taken using the optical sampling spatial oscilloscope from (a) the even and (b) the odd columns from all the nodes in the system at 160 Mb/s. Because the even and odd columns had different receivers, different voltages were applied in the two figures.

Before the 18 lasers were connected to the system, it operated for an extended time as a 2 x 2 video switch. In this mode of operation, video cameras were connected to codecs that encoded the video into ATM cells inside a SONET OC-3c frame at 155 Mb/s. The optical output from the

codecs was fed to a conversion board that converted the 1300 nm multimode signals to 850 nm single mode signals required by the switch. After passing through the switch, the detected outputs were converted back to 1300 nm multimode signals and sent to the codecs for conversion back to video, where the outputs were displayed on television monitors. This system was displayed at the National Communications Forum Tech Previews for several days and operated in the lab for several months after that. After transportation back from the show, no alignment was necessary. The system was subjected to rather harsh vibrations, including being shipped in tact in a truck and being dragged on rollers across the parking lot. This is not unusual for an optical system to be subjected to such vibrations and still operate, indeed robust optomechanical systems are routine for space and avionics applications. However, it is perhaps an important point that systems such as these can be built with robust mechanics, because in a qualitative sense, those in the microelectronics world are skeptical.

4.3.12 Multichannel optical oscilloscope for sampling optoelectronic circuits

The advent of large-scale, free-space, opto-electronic interconnections, as demonstrated in recent system prototypes [117,104,150,14], requires new sampling methods to reveal diagnostic information. Several factors contribute to the difficulty of probing optical communications channels without disrupting their operation. High-speed electronic connections to the chip periphery are not available in sufficient number and would contribute an undesirable thermal load. Electronic and optical [151,152] physical contact probes would obscure many of the optical channels that are relayed to a common surface of the chip in current systems. Optical sampling provides the better method although many standard techniques are either too time consuming or complex to implement.

We will describe a tool we developed that delivers diagnostic information on a large number of high-speed, optical data channels simultaneously and operates analogously to the conventional sampling electronic oscilloscope. The optical oscilloscope is constructed using CCD cameras and video capture boards that are controlled by a software application resident in a personal computer. Sampling is based on a stroboscopic method of using short pulsed laser probe beam synchronized to a data stream to illuminate optical modulators within the opto-electronic circuit. We have demonstrated and will discuss the tool's capability of simultaneously monitoring arrays of broadband opto-electronic devices operating at speeds from several hundred Megabit/s to a few Gigabit/s.

Fundamental Operation

In current free-space photonic systems, data is transmitted optically between electronic processor cells by modulating the intensity of light beams. Arrays of light beams, externally generated by laser diodes and diffractive components, are focused by lenses onto small reflective windows underlying multiple-quantum-well (MQW) material. The optical absorption of the MQWs is electronically governed by attached processing circuitry. In this manner, the absorption of the MQW encodes data onto each optical channel. An optical infrastructure then routes the reflected modulated channels to the subsequent chip or fiber.

An optoelectronic chip may embody thousands of optical channels each operating at speeds of hundreds of Mbits/sec, thus posing a serious challenge in collecting diagnostic information. In present-day investigations, it is typically necessary to simultaneously monitor a large number of parallel channels to determine the optimal operation parameters. Rather than design a complex

array of high-speed photodetectors that must be accurately aligned to a remote image of the modulator array, it is far simpler to sample the modulators' states using a repetitive, short duration light pulse and collect the image with an inexpensive CCD video camera. Thus in the same manner that a stroboscopic light source apparently freezes or slows the motion of rotating fan blades, the pulsed illuminator highlights the evolution of a periodic data stream for a large set of modulators.

Tool Design

A schematic of the multichannel, optical oscilloscope is shown in figure 60. The three primary

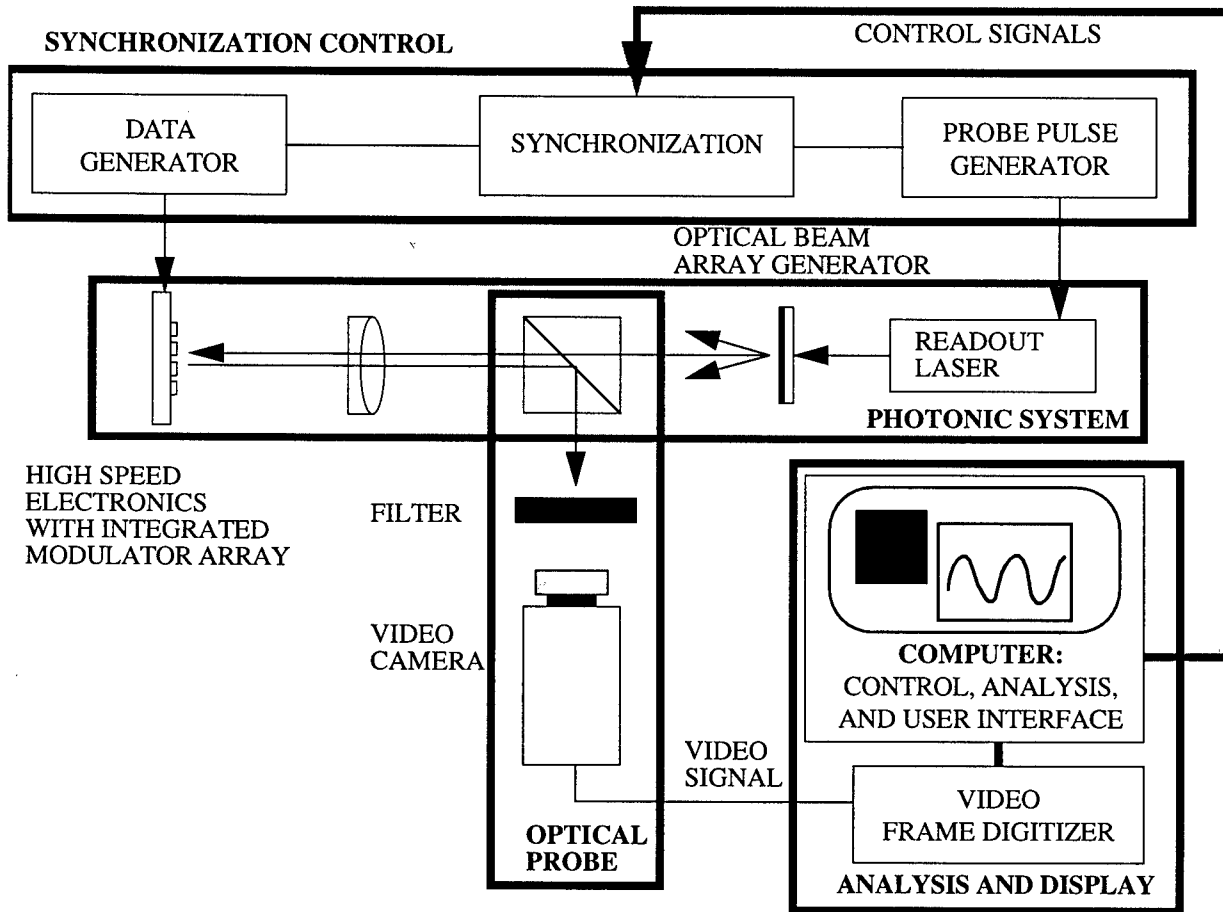


Figure 60. Schematic of test photonic system and optical oscilloscope modules.

functions of the oscilloscope hardware modules are:

- to generate pulses that create short duration readout light beam arrays for probing the modulator absorption and to synchronize these pulses with a periodic data stream (synchronization control module)
- to sample the readout light from several optical data channels in parallel and focus the individual channels separately onto a photosensor array, typically a CCD video camera (optical probe unit)

- to digitize and analyze the video signal and display the sampled waveforms in a format similar to that of an oscilloscope (analysis and display processor)

Currently, the separate modules have been only loosely integrated since the system to be investigated influences the design of the probe and synchronization units.

The synchronization control unit is typically custom designed to match the system. For example, generation of the short-duration readout light pulses can be performed by connecting new signal lines to the existing readout lasers in some systems, while in other cases a pulsed, broad-area illuminator can be integrated with the optical probe assembly. Since the readout pulse usually occurs repeatedly during the time sampling window of the photosensor, it must occur at the same point in the data stream throughout that window. In our demonstrations, we have maintained synchronization by using coupled data and pulse generators referenced to a common clock, but differing by about 1 Hz at the bit frequency so that the pulse slowly scans the data pattern. It is also possible to use a delay generator to scan the readout pulse through the data stream, although this would require control signals from the analysis unit to coordinate the process. The end result is that a multitude of samples are collected from a finite duration window of the data stream by a low speed photodetector.

The optical probe, also referred to as the viewport, extracts light from the system using a partially reflecting surface, such as a beam-splitter or pellicle. The deflector must be designed so that it does not seriously disturb the readout beams or other optical channels communicating information to the optoelectronic chip. The optical channels of systems we investigated operated at a wavelength of 850 nm which is within the sensitivity range of most CCD cameras. Filters are used to reduce the intensity level as necessary. Lenses form an appropriately sized image of the optoelectronic modulator array on the video camera. At this point, the user is able to watch the evolving intensity variation of the spots on a video monitor.

The analysis and display processor must digitize and store the video signal frame-by-frame. Enhanced multimedia computers are available with internal video cards that provide accessibility of the video memory to the processor. We have developed a custom written software application to coordinate the analysis, control, and display interface for the oscilloscope. Our implementation of the oscilloscope was produced for an Apple Macintosh 840AV and demonstrated on other compatible Macintosh platforms.

The location and sizes of the regions of interest (i.e., the modulator spots) are interactively defined by the user while examining either a live or captured video frame. During operation, the processor calculates the average intensity in a set of predetermined regions of interest and plots these intensities in a variety of possible waveform formats. Control of other display aspects, such as the scales of the intensity and time axis, are also provided to the user. In operation, our system collected and analyzed about 10 video frames/sec. The intensity resolution of the waveform is limited by video noise and the digitization accuracy of the A/D video signal convertor. The temporal resolution is effected by the width of readout pulse and the speed with which it scans through the data stream. The tool has proven capable of analyzing as many as 256 channels simultaneously without any degradation in performance.

One further advantage of collecting the data using a video camera is that the video signal can be recorded using video tape recorders. Thus, the performance can be archived and reanalyzed at a later time.




Figure 61. Readout light beams illuminating high-speed MQW modulator

High-speed system demonstration

To demonstrate the high-speed capability of the optical oscilloscope, an array of independent, electrically driven, differential multiple quantum well modulators [153] were monitored while operating at multi-Gbit/s speeds. The sixteen modulator windows are shown with their associated readout beams in figure 61.

The synchronization between the data stream and probe pulses were fixed by using two high-precision, frequency stabilized analog signal generators synchronized to a common clock to trigger digital data and pulse generators. The optical probe pulse was measured to have a width of about 200 ps as measured by an independent high-speed detector.

Two independent NRZ data waveforms and their complements were connected to four of the eight modulator pairs while a square waveform was connected to the remaining four pairs. Figure 62 shows the oscilloscope traces for data streams of 1 Gbits/s and square waves of 1GHz. All traces share a common time and intensity scale. The reduced intensity variation in certain channels is due to local heating, caused by termination of transmission lines, that shifts the operating characteristics of the modulator. Although the beams are well focused on the modulator windows in this




Figure 62. Optical oscilloscope traces from high-speed modulators for 1Gbit/s data and 1GHz square waves.

test, we have demonstrated that equivalent oscilloscope performance is attained when the spots are defocused to illuminate a larger area.

Free-space photonic switch fabric

To demonstrate its operation in a practical system, the multichannel optical oscilloscope was used to examine the performance of system 6. This demonstration system is composed of one optoelectronic chip comprising sixteen independent 16x16 crossbar switches, and the optomechanical infrastructure to relay optical channels between the chip and a two-dimensional fiber bundle array. The optoelectronic chip is a hybrid combination of GaAs MQW modulators bonded to VLSI silicon processing circuitry and illustrates the potential for dramatically expanding data throughput. The fiber bundle serves to collect and concentrate the data streams for remote external transmitters and receivers. In addition, a small number of low-speed electronic connections supply control and switch configuration information to the chip. During operation, the switching fabric has been shown to route digitized video and ATM-like traffic.

In this demonstration, one 16x16 switching node was examined. Two independent optical input channels provided periodic 8-bit data streams for this node. The switch was configured to route each of the two input streams to separate output modulators. The system was operated at a channel data rate of 200 Mbits/s which, under normal operation, provides sufficient overhead to allow switch reconfiguration and data encoding and a net communications channel throughput of 155 Mbits/s.

The photonic switch required only minor adjustments to enable the oscilloscope to monitor oper-

ations. A low reflectivity fused silica substrate, that was also used for inspecting optical beam registration, was inserted near the system pupil to deflect a portion of the modulated light toward a video camera. Lenses and attenuators were then selected to produce an image of the 16 output modulators residing on the surface of the optoelectronic chip. This viewport assembly did not disrupt system operation as evidenced by the undisturbed routing of a separate video data stream.

In order to coordinate the data and readout signals, a common reference clock synchronized a data generator and pulse generator operating at 200,000,001 Hz and 25 MHz respectively. The data generator, a Tektronix HFS9000 stimulus generator, is part of the system hardware. The readout laser, usually operated as a CW source, was modulated by the 25 MHz signal with a pulse width of about 1 ns. Thus a series of eight data bits were scanned over a time interval of about 8 seconds.

Figure 63 shows oscilloscope traces that were obtained simultaneously from the active system.




Figure 63. Fifteen traces obtained simultaneously from free-space photonic switch operating at a channel rate of 200 Mbits/sec. One of two separate input data channels is routed to each output modulator.

Only fifteen traces are displayed since only 15 of the 16 output modulators were illuminated by readout beams. This scheme is due to a design experiment where two alternate receiver designs, optimized for different power levels, were tested. The optimal configuration required shifting all spot arrays by one location so that the sparsely populated fiber array was aligned with the better performing circuitry leading to a more robust system demonstration.

During the diagnostics test, the multichannel oscilloscope provided the opportunity to investigate the full set of output channels that was inaccessible with the current, sparsely populated output fibers. In addition, it was possible to immediately discern the differing operating characteristics of the receiver designs as a function of the bias voltage.

Summary

Ongoing development of new opto-electronic interconnection architectures and technologies requires the invention of new diagnostic tools to investigate the performance of these novel photonic circuits. The application of stroboscope techniques using commercially available CCD cameras to probe optical absorption characteristics of modulator-based systems provides the basis for our high-speed, multi-channel, optical oscilloscope. We have demonstrated its capabilities by collecting diagnostics from parallel multi-Gbits/s data streams and from practical free-space photonic prototype systems.

4.3.13 Conclusion

We have described a new optoelectronic switching system demonstration that implements part of the distribution fabric for a large ATM switch. The system uses a new novel architecture, new device technology, a new optical system and new mechanical system. The system is a major step forward compared to our previous systems, primarily in that the device technology is made using high yield low cost manufacturable VLSI processes, the optical system is drastically simpler, and the opto-mechanical system is robust. The technology is now advanced enough that Terabit optoelectronic switching systems can now be contemplated within the next few years.

4.4 System 7

4.4.1 System 7: Goals and Objectives

System 6 was a successful demonstration in that it demonstrated concurrently:

- A large number of optical interconnections to a silicon integrated circuit
- No realignment for over a year
- Reliable operation at 155 and 208 Mb/s
- Simplified architecture which led to simplified optics
- No polarization components
- A single fiber array with both input and output fibers

There were several steps that needed to be taken, even after system 6, before this technology could be used in products. Those identified in the early stages were:

- Reduction in size so that the system could be placed on a PC board
- Increased data rate
- Manufacturable fiber array technology
- Operation with low cost off-the-shelf laser diodes for data inputs
- An application that absolutely required it.

System 7 set about to answer the first four points above.

A preliminary analysis and design of a compact optical system is given in section xxx. Initial estimates of its size indicate that it is roughly 1 x 1 x 3 inches compared to roughly 3 x 3 x 16 inches for system 6.

In the following section, we describe results from two switching chips capable of operating above 622 Mb/s.

We began a program to manufacture 2-D arrays of fibers, although this program was ended once the contract was finished.

More sensitive receivers without optical fan-out allowed low power lasers. The resultant system without a binary phase grating for this fan-out would be wavelength insensitive and require only ~ 250 microwatts of optical power.

4.4.2 System 7 Switching Chips

We describe initial results from two opto-electronic switching chips, one with 1024 differential optical inputs and 1024 differential optical outputs with individual channels tested above 600 Mb/s and a second with 512 differential optical inputs and outputs with individual channels tested up to 900 Mb/s. The technology for the chip consists of flip-chip bonding of 850 nm GaAs/AlGaAs multiple quantum well (MQW) detectors and modulators onto silicon CMOS with substrate removal to allow access to the optical devices [19].

The switching chip implements part of a simplified distribution fabric for a growable packet ATM switch [101,102]. Previously, an opto-electronic chip [73] and system [14] have been demonstrated implementing this architecture. That chip, designed using 1 μm CMOS, contained two hundred fifty-six 16 x 1 nodes (or sixteen 16 x 16 switches with optical fan-out) operating at a maximum speed of 450 Mb/s. The system operated at 208 Mb/s as a time multiplexed switch, capable of routing ATM cells at OC-3c rates (155 Mb/s) with an appropriate out-of-band controller. The current chips, designed in 0.8 μm and 0.35 μm CMOS, contain sixty-four and thirty two 16 x 16 switches respectively (Fig. 64.). The 16 x 16 switches are implemented by fanning out the *electrical* outputs from 16 differential receivers to sixteen 16 x 1 multiplexers, each with a differential optical output. The optical inputs and outputs are arranged in a rectangular array on top of the multiplexers. Control of the chips is electronic. The combination of increased density, the use of a third level metal with circuitry underneath the flip-chip bonding pads [111], and electrical fan-out allows four (0.8 μm) to ten (0.35 μm) times the functional circuitry in the same area. Table 16 gives a summary of the characteristics of the previous chip [73] and the two new chips.

	Chip 1 [4]	Chip 2 [7]	Chip 3
Technology	1.0 μm ($\lambda=0.5$)	0.8 μm ($\lambda=0.4$)	0.35 μm ($\lambda=0.24$)
Pixel function	16 x 1 node (mux)	16 x 16 switch	16 x 16 switch
Number of processing units	256	64	32
max data rate @ $< 10^{-10}$ BER	450 Mb/s	600 Mb/s	790 Mb/s *
best case sensitivity at design rate	-18 dBm @200Mb/s	-15dBm@600 Mb/s	-15dBm@625 Mb/s
required optical energy/beam	(~80fJ)	(~50 fJ)	(~50 fJ)
Optical I/O	4096 in 256 out	1024 x 2 in 1024 x 2 out	512 x 2 in 512 x 2 out
Potential Throughput	115 Gb/s	600 Gb/s	404Gb/s
Potential I/O Bandwidth	230 Gb/s	1.2 Tb/s	800 Gb/s
Tested Throughput	2.92 Gb/s [5]	600 Mb/s	790 Mb/s
Chip size	7 x 7 mm	7 x 7 mm	4 x 4 mm
Optical field size	5.44 x 5.44 mm	5.12 x 5.44 mm	2.24 x 2.24 mm
Optical window spacing	80 μm	80 μm	35 μm , 70 μm
Window density	15.6K/cm ²	15.6K/cm ²	40K/cm ²
Electrical static dissipation	1W @ 5.5 V	5W @ 5V	0.3W @3V
#FETs	160K	450K	225K
Electrical I/O	23 @ 120 Mb/s	23	23

Table 16. Characteristics of previous chip and the two new chips. (*limited by test equipment)

The receiver for chip 2 (Fig. 65a.) uses a modified design of that described in [21]. The sensitivity versus bit-rate (Fig. 65b.) was measured for a single receiver. The performance is not quite as

good as that in [21], one reason for this, is that an imbalance was intentionally introduced in the stage following the receiver before a large electrical fan-out driver to reduce the static dissipation. A second reason is that the receiver outputs from neighboring 16×16 switches were inadvertently pair-wise connected and testing was done by driving both receivers in tandem using a 1×2 binary phase grating. Unequal powers on the receivers may cause signal distortion in the common line. Our previous chip with single ended receivers showed good uniformity ($< \pm 400$ ps gate delay variations) across the chip [73]. We expect the variations in these circuits to be comparable or less.




Figure 64. (a) Block diagram of chip. Each square in represents a 16×16 switch. (b) Each switch is implemented using sixteen 16×1 multiplexers with electrical fan-out of receiver outputs (arrows). (c) optical differential receivers (light color) and transmitters (dark color) that are overlaid on top of the switches. Receivers are mapped by number to horizontal fan-out lines (arrows) and transmitters are mapped from multiplexers (column numbers)

The large number of active receivers leads to a static power dissipation of almost 5W. The exciton shift due to thermal heating of the circuit was found to be ~ 10 nm both at the center and at the edges of the array. Using a shift of 0.28 nm/ $^{\circ}\text{C}$, the thermal resistance of the package was found to be 7 $^{\circ}\text{C}/\text{W}$. The uniformity of the position of the exciton peak in reflectivity while dissipating 5W indicates the temperature uniformity across the array is within a few degrees. The chip mount was specifically designed for temperature uniformity [86].




Figure 65. (a) receiver schematic and (b) sensitivity versus bit-rate for the switching nodes in chip 2.

There are 16,384 ($64 \times 16 \times 16$) paths through the switching chip. Sixteen paths were measured at a data rate of 625 Mb/s (Fig. 66), sampling the upper left 16×16 switch, exercising each

receiver and each modulator driver from that switch. Due to the architecture of the chip, we would not expect large delay variations between inputs, this was indeed the case on our earlier chip that was more thoroughly characterized [73]. Care was not taken to ensure uniformity of photocurrent throughout this set of measurements, so there might be some power dependent variations in delay. Individual channels had bit error rates below 10^{-11} at 600 Mb/s and below 10^{-9} at 625 Mb/s. While all paths weren't measured, all but 7 MQW diodes luminesced under forward bias. Thus, if there are no additional problems in the silicon circuit, nearly all paths should be operational.




Figure 66. Sixteen outputs from the upper left 16 x 16 switch with output i selecting input i at 625 Mb/s with a bit pattern of 01000111 for chip 2. The zero line is accurate, i. e. the contrast ratio is approximately 2:1.

The chip was redesigned in 0.35 μm CMOS. Because of the decreased feature sizes, another factor of 2.5 in density was achieved. The receiver design eliminated the second stage gain broadening transistors and included clamps in the feedback path. Elimination of the gain broadening transistors reduces the static power dissipation at the expense of increased sensitivity. If the feedback FET is gated off, the clamps further reduce power dissipation and allow dynamic operation, which has not yet been tested on this chip.

This chip had a few minor design errors, which are currently being corrected. We were still able to pass signals through many paths of the chip and control many of the nodes. Figure 67 shows operation of a single channel at 800 Mb/s and 900 Mb/s with pseudorandom bit patterns. Bit error rates were measured up to 790 Mb/s which was the maximum rate that the BER receiver would respond to a direct laser input. (The specification for the BER receiver was 700 Mb/s). We might expect operation above 900 Mb/s judging by the eye diagram and this will be investigated further.

We have described initial characteristics of opto-electronic switching chips with poten-

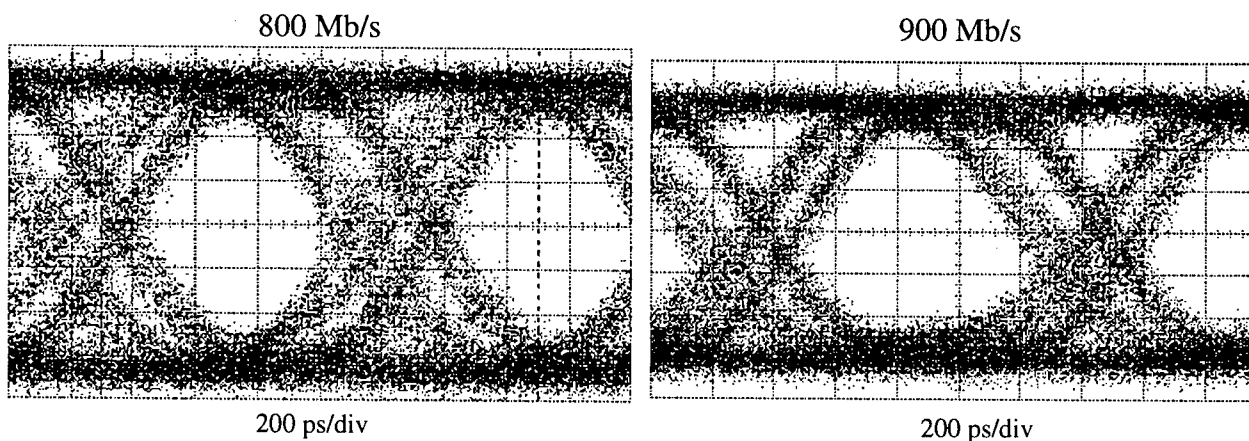


Figure 67. Operation of the 0.35 μm switching chip at 800 and 900 Mb/s. Optical power was roughly 250 μW per beam.

tially greater than 1 Terabit per second I/O bandwidth. While crosstalk, thermal spatial variations under dynamic operation, and delay variations need to be rigorously characterized before this I/O bandwidth can be utilized in a system, these chips further illustrate the potential of hybrid opto-electronic VLSI smart pixel technologies.

4.5 Future Optical Systems

4.5.1 Introduction

As we mentioned in the discussion of system 7, reduced physical size along with reduced cost were important attributes that need to be addressed. In this section, we describe two aspects of this work. First we describe a general analysis of these types of systems. It is this understanding that led to the compact designs that we hoped to use in system 7. In the following section, we describe a class of all-diffractive optical interconnects. The diffractive design overcomes the problem of chromatic aberration that accompanies many diffractive designs, yet can access larger arrays of devices separated by longer distances than a traditional "micro-channel" approach.

4.5.2 Paraxial Optics of Free-Space Photonic Switching Systems

Introduction

Wide-angle lenses and compact arrays of fibers and modulators can be used to design compact optics modules for photonic switching systems. The tradeoffs between aggregate I/O data rate and optical system complexity can be quantified using the optical invariant. The use of a common optical design tool, the y - \bar{y} diagram, in photonic switching system design is demonstrated. This design tool is used to find specific relationships 1) between system parameters and optics module size and between, and 2) between image properties and aggregate system I/O bandwidth.

Over the past several years, tremendous advances have been made in the field of free-space photonic switching [14,104,117,118]. Many recent efforts in photonic switching use optics modules that are similar to one another. These optics modules contain an input, a switching device, and some sort of optical fanout. The input is an array of spots in a telecentric object plane. The switching is performed by an array modulators and detectors in a telecentric image plane, and optical fanout is performed in a collimated pupil plane.

Since many of photonic switching system demonstrations use similar optics modules[154], it is worthwhile to closely examine the fundamentals of this type of optical design. Designing free-space photonic switching systems requires a wide range of expertise, including CMOS design, electronics design, optoelectronic design, optomechanical design, lens design, and system architecture. Often, a design choice in one discipline may seem to be an obvious choice, but may have profound implications on the other aspects of the system design. Better systems can be built if the tradeoffs among the various disciplines are better-understood across the disciplines. The fundamentals of the optical design have serious implications on the size of the system and the tradeoffs between aggregate I/O bandwidth and optics module complexity. This paper describes some of the fundamental limitations of optical design and the implications on issues such as system complexity and aggregate I/O bandwidth.

Photonic switching systems have goals of high aggregate throughput as well as simple, robust, and compact physical construction. Unfortunately, these goals often conflict; for example, maintaining high aggregate throughput may force the system to have small detectors, which forces the lenses to have a low f -number, which, in turn, forces the optics module to be complex. Many tradeoffs are required between system size, system complexity, and system performance.

This paper describes some of these tradeoffs and the optical principles underlying the tradeoffs.

This paper presents several pieces of information that are useful to the photonics switching community. First, the paper contains an extended example of the use of the y - \bar{y} diagram for photonic switching systems. Second, it presents in Section III a concrete way to examine how fiber bundle pitch, detector pitch, number of channels, and field of view of the lenses affect the overall system length, as well as a detailed discussion on possibilities for optimizing each of these variables. Third, it presents in Section VI a concrete way to examine how optical system complexity, wavelength, spot size, detector pitch, and receiver design combine to affect aggregate system I/O.

Paraxial properties of digital photonic switching systems

Similar optical designs are used in many different photonic switching systems. These systems have an input array of spots in the object plane, an array of detectors and modulators in an image plane, and an output array of spots in another image plane. Lenses are used to transfer the input array of spots onto the array of detectors, then onto the output array.

The input array of spots can be an array of optical fibers, an array of vertical cavity surface emitting lasers (VCSELs), or an array of modulators. Each of these input arrays can force the object plane of the system to be telecentric.

The functional fanout of the system is often performed optically, using a diffraction grating. Vignetting by the objective lens is eliminated by placing the grating in aperture stop. Separate paths for the input and output beams are provided by using pupil division.

Processing and switching in these systems is generally performed in an image plane, using reflective devices such as optoelectronic VLSI chips (OE-VLSI)[73]. To be able to get information both onto and off of the device, it is often placed in a telecentric image plane.

Excellent insight into optical system design can be obtained using a tool called the y - \bar{y} diagram. A y - \bar{y} diagram is a parametric plot of the height of the marginal ray (y) in an optical system as a function of the height of the chief ray (\bar{y}), with axial position in the system as the parametric variable. On the y - \bar{y} diagram straight y - \bar{y} lines represent the spaces between lenses, and bends in the y - \bar{y} line represent lenses (or mirrors with power.)

The simplest possible optical system with the required paraxial properties for photonic switching, telecentric object and image planes and a collimated pupil plane, contains only two surfaces with power. The y - \bar{y} diagram and system schematic of such a system are shown in Figure 68.

Although a simple two-lens system such as the one represented in Figure 68 would not have the required image quality, this kind of simple system is useful as a comparison for other, more realistic systems. For example, any system shorter than this simple, two-lens system can be considered a compact system.

A y - \bar{y} diagram such as the one shown in Figure 68 is an extremely efficient method for representing the paraxial properties of an optical system[155,156,157,158]. The leftmost vertical




Figure 68. Two paraxial representations of the design of the same photonic switching system. Both representations show a telecentric object space with object height \bar{y}_o , collimated aperture stop of radius y_a , and a telecentric image space with image height \bar{y}_i . Figure 68 a) is a y - \bar{y} diagram. Figure 68 b) is a system schematic. Although the system schematic shows information about ray angles as well as ray heights, the y - \bar{y} diagram often provides more insight to the optical designer.

line in the figure represents object space; the diagram shows both that object space is telecentric (because the line is vertical at the intercept) and that the object height is \bar{y}_o (the value of the intercept with the \bar{y} -axis.) The horizontal line represents the aperture stop space; the line shows both that the aperture stop is in collimated space (the line is horizontal at the intercept) and that the aperture stop has a half-diameter of y_a (the value of the intercept with the y axis). Similarly, the right-most vertical line represents image space; the diagram shows both that image space is telecentric (because the line is vertical at the intercept) and that the image height is \bar{y}_i (the value of the intercept with the \bar{y} -axis.)

Similar, though less general, information can be obtained from a simplified sketch of the imaging optics of a photonic switching system, as shown in Figure 68b. This figure also shows that the object and image spaces are telecentric (the chief ray slope is zero) and that the aperture stop space is collimated (the marginal ray slope is zero.) Furthermore, this figure also shows the

size of the object and image (the chief ray heights when the marginal ray heights are zero) and the size of the aperture stop (the marginal ray height when the chief ray height is zero). Additional information from this figure includes the length of the system, the numerical aperture in object and image space, and the field of view of the lenses.

Although system schematics such as the one in Figure 68b show more information about a completed system design, y - \bar{y} diagrams such as the one in Figure 68a are often more useful to the optical designer. The main advantage of the y - \bar{y} diagram is that it deals directly with the required image and pupil sizes and locations and determines the required lens properties afterwards; other methods, on the other hand, deal directly with the systems' lens properties, obtaining the image qualities only as a secondary calculation. A second reason y - \bar{y} diagrams can be more useful to the optical designer is that they are more general; they don't require the designer to establish a numerical aperture of the system, for example. A third reason y - \bar{y} diagrams can be more useful than system sketches is that the requirements of changes in conjugates or changes in system length can be determined naturally and easily with the y - \bar{y} diagram.

Length in a system is proportional to area on the y - \bar{y} diagram. For the system shown in Figure 68, the length from the object to the image is given by:

$$l = 4 \cdot \frac{A}{W} = \frac{4 \cdot (\bar{y}_o + \bar{y}_i) \cdot y_a}{W} = 2 \cdot \frac{(\bar{y}_o + \bar{y}_i)}{\tan\left(\frac{1}{2}FOV\right)} \quad 1$$

where l is the length between the object and image planes, A is the area on the y - \bar{y} diagram swept out by the y - \bar{y} line, W is the optical invariant, y_a is half the entrance pupil diameter of the lenses, \bar{y}_o is the height of the chief ray at the object (half the diagonal of the fiber array), and \bar{y}_i is the height of the chief ray at the image (half the diagonal of the OE-VLSI). The optical invariant is a fundamental measure of an optical system and in air is given by

$$W = \bar{y}_k \cdot u_k - y_k \cdot \bar{u}_k \quad 2$$

where y_k represents marginal ray height at plane k , u_k represents marginal ray slope at plane k , \bar{y}_k represents chief ray height at plane k , and \bar{u}_k represents chief ray slope at plane k . The optical invariant has the same value for every plane in the optical system[159].

The length of a real photonic switching system differs from the paraxial system described in

Equation 1 by a scaling factor α :

$$l = 2 \cdot \alpha \cdot \frac{(\bar{y}_o + \bar{y}_i)}{\tan\left(\frac{1}{2}FOV\right)} \quad 3$$

The scaling factor α is a measure of how compact an optical design is with respect to the idealized system represented in Figure 68. Scaling factors for unusually compact systems such as cassegrain telescopes can be as small as 0.15; scaling factors for unusually large systems such as high-performance zoom lenses can be as large as 10. Scaling factors for lenses used in photonic switching systems are often fairly modest ($1.5 < \alpha < 2.5$), making the systems slightly larger than the corresponding paraxial system.

Reducing system size

To be competitive with electronic switching systems, photonic switching systems must become smaller and more robust. One plausible goal for the optics modules in photonic switching systems is to mount them on the electronic circuit boards which will be part of such a system. This section describes some tradeoffs that can be made towards this goal.

Equation 3 shows that there are several mathematical possibilities for decreasing system size (l):

- maximize the field of view of the lenses (maximize FOV)
- minimize the size of the OE-VLSI array (minimize \bar{y}_i)
- minimize the size of the fiber array (minimize \bar{y}_o)
- make the system shorter than its paraxial counterpart (reduce α)

The most promising method for reducing the size of photonic switching systems is to maximize the field of view of the lenses. Current systems use a 1/2 field of view of only 5-7 degrees; this value may be increased to over 35 degrees. Exactly how far the field of view can be increased is not clear. Some lenses, such as endoscopes and door viewers have fields of view over ± 80 degrees[160,161]. However, these systems are unacceptable models because they have unacceptable distortion, do not have telecentric image planes, and have an internal aperture stop. Wide-angle scan lenses, on the other hand, provide excellent starting points because these lenses, like the lenses required for photonic switching systems, have an external aperture stop in collimated space, have carefully controlled distortion, are typically designed for only a single wavelength, and are typically diffraction-limited. A wide scan angle for these lenses is ± 40 deg[162]. Unfortunately, these systems are not at all compact; they are often as much as 8 times as long as their focal lengths. A reasonable field of view for a compact lens in photonic switching systems may be ± 30 deg.

The next most likely method for reducing the size of photonic switching systems is to

minimize the size of the OE-VLSI (reduce \bar{y}_i) Image size in a current system demonstrator is $\bar{y}_i = 3.62\text{mm}$. The size of the image for a square device array is determined by the number of modulators and the pitch of these modulators.

$$\bar{y}_{img} = P_{img} \sqrt{\frac{N}{2}} \quad 4$$

where P_{img} is the pitch of the modulators and N is the number of modulators. The number of modulators is usually determined by factors such as the required functionality of the switch and the modulator fabrication capabilities. A reasonable number of modulators is $N=64 \times 64$. The pitch of the modulators is determined by factors such as the window size, the size of the solder bumps, and the space required for processing circuitry. The pitch for OE-VLSI devices in a current system demonstrator is 0.080 mm. Although window size places a lower limit on the window spacing, the size of the solder bumps, which are used to bond the GaAs-AlGaAs detectors and modulators to the Silicon CMOS circuitry, usually is a more important consideration. The solder bumps are used as bond pads to bond the GaAs detector and modulator windows onto the Silicon CMOS circuitry. Chips with pitches as small as 0.035mm x 0.070mm have been fabricated. For future planning, a reasonable assumption might be that the total area required for each window can be reduced by a factor of 3, yielding a detector pitch of about 0.030mm for a square array. This spacings implies an image size of $\bar{y}_i=1.36\text{mm}$ for an array of 64x64 windows.

The final method for reducing the size of photonic switching systems is to minimize the size of the object (reduce \bar{y}_o). For systems with a fiber bundle for an object, the size of the object is determined by the size of the array, the diameter of the fibers, and the required spacing between the fibers. The values for these parameters in a current system demonstrator are: a 64x64 array, 0.125mm diameter fibers, and 0.125mm between adjacent fibers, giving a fiber bundle size of about $\bar{y}_{obj}=11.3\text{mm}$. The size of a square object is

$$\bar{y}_{obj} = P_{obj} \sqrt{\frac{N}{2}} \quad 5$$

where P_{obj} is the pitch of the fibers. The pitch of the fibers is determined by the size of the fibers and the space required between the fibers. Current pitches are about 0.250mm, using 0.125mm diameter fibers with 0.125mm separation[73]. Extreme measures have yielded linear arrays of fibers with pitch as small as 0.140mm. Extending this spacing to a two-dimensional array would yield an object size of $\bar{y}_o = 6.33\text{mm}$ for an array of $N=64 \times 64$.

Making the system shorter than the corresponding paraxial system is mathematically possible, but is probably not technologically possible. The y - \bar{y} diagram for a short system is shown in Figure 69. The optical system represented in this figure shorter by about a factor of two ($\alpha = 0.5$)

as compared to the simple system represented in Figure 68, assuming identical optical invariants; this reduction in system length is evident from the reduced area swept out by the $y-\bar{y}$ line. The optical system represented in Figure 69 differs from the simple system represented in Figure 68 in three respects. First, positive field lenses have been added near the object and image; this can be seen in the figure by the clockwise bend in the $y-\bar{y}$ line at the points marked a) and e) in the figure. Second, a negative lens has been added between the object and stop; this addition can be seen in the figure by the counter-clockwise bend in the $y-\bar{y}$ line at point b) in the figure. Third, positive lenses have been added around the stop; this can be seen in the figure by the clockwise bend in the $y-\bar{y}$ line at the points marked c) and d) in the figure. This reduction in system length is achieved without changing the object, image, or stop properties; this can be seen by noticing that the $y-\bar{y}$ has the same intercepts with the y -axis and the \bar{y} -axis in both Figure 68 and Figure 69. Unfortunately, the design of this compact system would be technologically daunting; one of the unfortunate lessons of lens design is that the length of a system often determines the potential of its image quality more than the number of surfaces or the ingenuity of the designer[163]. Therefore, any buildable, durable system will probably be significantly longer than the one shown in Figure 68. A reasonable future system would have a length slightly longer than the one in Figure 68 ($\alpha = 1.2$). Moving to extreme methods in optical design may reduce the system length to be equal to the one shown in Figure 68 ($\alpha = 1$)




Figure 69. A $y-\bar{y}$ diagram for a size-reduced photonic switching system. Like the system represented in Figure 68, this system has telecentric object space with object height \bar{y}_o , collimated entrance pupil of radius y_a , and a telecentric image space with image height \bar{y}_i . However, this system would be about half as long ($\alpha = 0.5$). The reduction in size is possible because of the addition of three extra lenses. However, this system would probably not be able to be designed with acceptable image quality.

Table 17. Comparison of current and future system lengths and an example system. Future digital photonic switching systems will be much smaller than current systems because of smaller input and output arrays and because of wide field-of-view lenses. Although the example system doesn't take full advantage of the possible reductions in system size, it demonstrates the feasibility of small systems.

	Current	Reasonable future	Optimistic future	Example system
array size (N)	64x64	64x64	64x64	64x64
detector pitch (P_{obj})	0.080mm	0.060mm	0.030mm	0.080mm
fiber pitch (P_{obj})	0.250mm	0.250mm	0.140mm	0.125mm
1/2 FOV	6.4 deg.	25 deg.	30 deg.	20 deg.
scale factor (α)	1.1	1.2	1	1.5
system length	392mm	72mm	27mm	76mm

The overall system length can be determined by substituting the equations for image height (Equation 4) and object height (Equation 5) into the equation for system length (Equation 3):

$$l = \sqrt{2N} \cdot \alpha \cdot \frac{(P_{img} + P_{obj})}{\tan\left(\frac{1}{2}FOV\right)} \quad 6$$

Combining the possibilities of making the system smaller than its paraxial counterpart, minimizing the size of the fiber array, minimizing the size of the OE-VLSI, and maximizing the field of view of the lenses, and substituting into Equation 6 yields opportunities for reducing the size of the optics of a photonic switching system. In the most optimistic case, yields a system length of 27mm; for the more reasonable cases, Equation 6 yields a system length of 72mm. The y - \bar{y} diagrams of the four systems are shown in Figure . The figure graphically demonstrates two important points. First, initial improvements in system size will be from increased field of view lenses; this point is evident from reduced intercept of the y - \bar{y} line with the y -axis between lines a) and line b). Second, the figure shows that the size of future systems will be limited by the size of the fiber bundles; this point is evident from the relatively large value of the intercept of the y - \bar{y} line

with the negative \bar{y} axis.




Figure 70. Schematic of a compact optics module for photonic switching.

Figure 70 shows an example of an optics module for a small optoelectronic switching system. This system contains several challenging fabrication aspects: some small edge thicknesses, some small center thicknesses, some thick meniscus elements, and an aspheric correction plate incorporated into the diffraction grating. Nevertheless, the system is diffraction-limited, has reasonable tolerances, and would work well with self-centering lens mounts.




Figure 71. y - \bar{y} diagrams of the paraxial designs of a) a current photonic switching system b) a reasonable future system, and c) a future system with optimistic assumptions on system parameters.

Packaging Issues

Optics modules of the size listed in Table 17 can be efficiently packaged in stainless steel barrels and sub-cell lens mounting, similar to high-quality microscope objectives[164]. However, constructing these tiny systems will require new challenges in the optical design of photonic switching systems. Currently, most systems are designed, constructed, and aligned as if the system consisted of two separate subsystems - a collimating lens and a focusing lens. To achieve the small system sizes listed in Table 17, the systems must be conceived and mounted as an entire system. Some developments toward this goal include: self-centering mounts for the fiber array, self-centering mounts for the modulator array, focus adjustments within lens sub-cells, and testing the system as an afocal system.

Alignment provides special challenges in the design of photonic switching systems. The optics module must not only form excellent images, it must also precisely align the array of spots onto the array of detectors and modulators. Systems as small as those listed in Table 17 might not have enough space for viewports to visually align the spots to the windows, so new alignment techniques must be developed. These new techniques may involve temporarily placing a camera in the image plane or examining the signals on the detectors using instruments such as an optical oscilloscope, which can simultaneously examine the modulation across a large section of an array[108].

Reducing the size of the optics module in photonic switching systems will tighten the tilt and decentration tolerances on the lenses. However, it is safe to assume that small systems such as these can be built since they are similar to microscope objectives in size and image quality. The tolerances on lateral object-image misalignments will be the same as on current systems since the window size will be essentially the same for the small systems. The rotational tolerances will be looser than in larger systems since the overall image size will be smaller.

Tradeoffs between aggregate I/O bandwidth and complexity

One of the essential trade-offs in the design of photonic switching systems is the tradeoff between window size and system complexity. Small windows are desirable because they offer a much higher signal bandwidth due to reduced capacitance. However, small windows require optics modules with low f-numbers, which are very complex. Many factors can make an optical system complex: many elements, unusual glass types, aspheric surfaces, and tight tolerances. Complex optical systems are expensive, difficult to design and build, and susceptible to environmental instabilities such as vibrations and temperature fluctuations. This section presents the tradeoffs between optical system complexity and aggregate I/O bandwidth, then demonstrates a method of quantifying the tradeoffs.

Although complexity of a single optical system is difficult to quantify, it is possible to make some general statements about the relative complexities of optical systems. For a given image quality, f-number, and system constraints, a system with a large field of view will generally need to be more complex than a system with a small field of view. Similarly, for a given image quality, field of view and system constraints, a system with a low f-number will generally need to be more complex than a system with a high f-number. There is also a rough correlation between the optical invariant and the difficulty of designing a system. Undergraduate students in their first lens design class can design a lens with an optical invariant of $W=0.1$ mm as a long homework assignment. A good optical design engineer can design a lens with an optical invariant of $W=1$ mm as a reasonable design task. Teams of full-time lens designers take several months to design photolithography systems with optical invariants of $W=4$ mm.

The optical invariant is a very fundamental measure of an optical system. It is closely related to conservation of energy, the space-bandwidth product. Because it is such a fundamental quantity, the optical invariant is also an imperfect measure of the required complexity of an optical system. Special requirements on distortion, chromatic aberration, system size, image size, system durability, or other factors can cause systems with similar optical invariants to vary widely in complexity. Furthermore, skill of the optical designer, skill of the optical technician, and even

luck can result in simple systems with large optical invariants.

These ideas can provide a feeling for the required complexity of the optics module in a photonic switching system. Since photonic switching systems often have similar requirements on image quality, distortion, image size, and environmental and mechanical stability, systems with smaller f-numbers have to be more complex than systems with large f-numbers, all other factors being equal. Similarly, a system with a large image field has to be more complex than a system with a small image field, all other factors being equal.

The comparison of system complexities for large vs. small fields and large vs. small f-numbers can be combined by comparing the systems' optical invariants. At the image plane, the marginal ray height is zero ($y=0$), so the optical invariant simplifies to:

$$W = \bar{y}_{img} u_{img} \quad 7$$

where \bar{y}_{img} is half the diagonal of the image surface and u_{img} is the numerical aperture of the image plane. In photonic switching, systems, the height of the image is equal to the half-diagonal of the detector array, given in Equation 4.

In photonic switching systems, the numerical aperture of the lens is determined by the window size:

$$u = S \cdot \frac{\lambda}{d} \quad 9$$

where λ is the wavelength ($\lambda = 852\text{nm}$ is often used), d is the detector size ($d = 0.010\text{mm}$ is a current value), and S is a scaling factor that determines the relative size of the spot and the window. ($S=1.5$ is common). Small values of the scaling factor S imply a spot larger than the window; large values of S imply a spot smaller than the window. Figure shows the relative beam size for several values of this scaling factor.

Optical trade-offs with window size

Detector and receiver bandwidth is often limited by capacitance of the window; in this limit, the bandwidth is inversely proportional to the window area. The aggregate system bandwidth B is the channel bandwidth times the number of channels N ,

$$B = N \cdot \frac{k}{d^2} \quad 10$$

where k is a constant of proportionality that is determined by factors such as CMOS line size, and

d is the linear dimension of the windows.

Substituting the equation for the image height (Equation 8), the equation for the numerical aperture (Equation 9), and the equation for the channel bandwidth into the equation for the optical invariant (Equation 7) yields:

$$W = \lambda \cdot S \cdot P_{img} \sqrt{\frac{B}{2k}} \quad 11$$

This equation provides a useful way for designers to make quantitative tradeoffs between optical system parameters and device parameters.

Equation 11 shows that increasing the aggregate system bandwidth B can theoretically be obtained by:

- decreasing the pitch (P_{img})
- improve receiver design (increase k).
- decreasing the scale factor (S)
- increasing the optical invariant (W)
- decreasing the wavelength (λ)

The most likely way to improve aggregate system bandwidth is to improve receiver design (increase k). This improvement in receiver design will proceed naturally as CMOS technology moves to smaller line sizes.

A second likely way to increase aggregate system bandwidth without increasing optics module complexity is to reduce the pitch (P_{img}). Although decreasing the pitch allows for a higher bandwidth, it also makes chip design and fabrication more difficult. A reasonable window spacing might be $P_{img}=0.030\text{mm}$ instead of the current $P_{img}=0.080\text{mm}$.

The next most likely possibility to increase aggregate system bandwidth without increasing optics module complexity is to decrease the scale factor (S). This decrease in scale factor can would probably be accomplished by decreasing the window size while keeping the spot size the same. In the past, the performance of many digital optical systems were limited by laser power. New design techniques, higher-power lasers, and more sensitive receivers have reduced this problem. In some future systems, it may be desirable to allow the spots formed by the optics module to be much larger than the windows. An added benefit of increasing the spot size relative to window size is reduced sensitivity to misalignments. Naturally, the pitch of the detectors places an upper limit on spot size; if the spots become too large, crosstalk between channels will become a problem.

Increasing the optical invariant (W) is not a very likely way to increase aggregate system bandwidth. The optical invariant of current systems is about $W=1.2 \text{ mm}$. Designing optics modules of this level of complexity is a challenging task for a good optical designer. The most likely

way to increase the optical invariant of a system without placing unreasonable demands on the optical designer is to release some constraints on the optics modules. These constraints include system size, ruggedness, distortion, and telecentricity. Unfortunately, it is very likely that other systems issues are likely to tighten these constraints and require the optical invariant to decrease.

Decreasing the wavelength (λ) is also not a very likely opportunity for increasing aggregate system bandwidth. Excellent lasers are available in the near IR. Also, excellent detectors and modulators are available in this region. Furthermore, 1.3 μ m and 1.5 μ m are far more common in communications systems, so it is likely that systems will move to these longer wavelengths rather than shorter ones.

Conclusions

The paraxial optical design of photonic switching systems can be readily made very small. Systems this size could be mounted on a circuit board, making them much more suitable to be included in electronic switching systems. This size reduction is possible with conventional optical fabrication technology, although the results presented here are general, and applicable to technologies such as diffractive optics, and GRIN lenses.

Fiber bundles are currently a limiting factor in the paraxial design of the optics modules for photonic switching systems. Fiber bundles are often used as an input plane to photonic switching. The size of the optical fibers and the limitations on how closely the fibers can be spaced places limits on the size of the fiber bundle, which, in turn, places limits on the overall size of the optics module. Advances in fiber-bundle technology or other technologies for input arrays present excellent opportunities for reducing the size of photonic switching systems.

Designing small photonic switching systems requires compact, wide-angle lenses with external stops. Laser scan lenses provide reasonable design types and starting points. However, using these design types in photonic switching systems requires significant modifications to make the lenses more compact, to give them the proper distortion properties, and to make them telecentric.

Reducing the pitch of the windows on the chip will increase the per-channel bandwidth of the system as well as decrease the overall system size. Reducing the pitch will place challenges on chip designers and chip fabricators. It is hoped that this paper will help these designers and chip fabricators to understand the tradeoffs involved in this endeavor.

4.5.3 Point to point interconnections using diffractive optical systems

To commercialize optoelectronic VLSI, researchers developing packaging technology must find a strategy that is both suitable to deliver the performance required and acceptable to conventional electronics system designers who are unfamiliar and perhaps uncomfortable with conventional optics. A short list of desirable system attributes would include: densely packed optical channels, inexpensive and compact packaging, and optomechanical systems without critical micron level alignment tolerances.

One approach is to define individual microchannels connecting each transmitter and receiver pair using sets of microlenses[165]. This scheme provides great flexibility in defining space variant interconnections. It also significantly reduces the lateral extent of the optics when compared to

most refractive approaches. Unfortunately, densely packed channels accentuate the gaussian nature of the beams thereby limiting the practical interconnection range. This is unsuitable for conventional circuit board layouts requiring connection lengths of about 10 to 200mm.

The interconnection range can be increased by using an intermediate relay lens shared by a set of optical channels[166]. In this paper we describe an all-diffractive-optical-relay (ADORE) where the microlenses and the relay lens are diffractive elements. The choice of diffractive lenses allows the use of lens arrays to partition a system, thereby extending the surface area that can be served. Diffractive lenses also contribute fewer aberrations than similar refractive lenses. In addition, when the alignment sensitive microlenses are aligned to the sources and receivers via accurate bonding techniques, the remaining assembly requires notably relaxed alignment tolerances. Finally, similar to other balanced grating systems[167], the chromatic sensitivity is reduced leading to an extended operating range.

The basic configuration of the ADORE optical elements is shown in figure 72. The two dimensional array of microlenses is partitioned into sets associated with a relay lens that is larger in diameter than the gaussian beam waist. Each lens in the microlens array has a slight lateral offset so that the beam is directed toward and forms a geometric focus at the center of the relay lens. The microlens arrays are positioned with respect to the relay lens to form a $4F_{\text{relay}}$ imaging setup. Both microlens arrays would typically be designed with the same focal length, $f_{\mu\text{lens}}$. It is also possible to incorporate magnification, leading to configurations that can generate a perfect shuffle interconnection. When studying the performance of the ADORE system, it is crucial to include the gaussian beam behavior since practical channel spacings are projected to be in the range from 62.5 to $250\mu\text{m}$.




Figure 72. Basic ADORE configuration. The top trace displays the geometric performance of the ADORE setup. The bottom trace illustrates the behavior of the gaussian beam waist. This scheme supports bidirectional interconnections.

Two critical issues must be addressed for a diffractive optics scheme: the chromatic sensitivity of the system, and, the influence of scattered light due to the limited diffraction efficiency. The chromatic sensitivity is typically a severe problem for diffractive lenses whose focal lengths vary inverse proportionally to wavelength. A limited chromatic tolerance would lead to strict VCSEL spectral requirements resulting in increased component price.

The problem of scattered light is primarily related to limitations of the fabrication process. The minimum feature size (about 1-2 μm) restricts the number of phase levels for faster lenses. For the moment, we will disregard the losses noting only that diffraction losses will be distributed fairly uniformly over a large area and that advances in direct e-beam lithography and gray-scale masks will improve microlens diffraction efficiency.

The ADORE configuration displays an astonishing low chromatic sensitivity. Figure 73 shows the simulated image size variation by wavelength for a range of focal length ratios, $F_{\text{relay}}/f_{\text{microlens}}$, for the case $f_{\text{microlens}}=345\mu\text{m}$. The lighter colored regions represent optimal conditions for operating the system. For larger $f_{\text{microlens}}$, the area of reduced sensitivity decreases. From the diagram, we can determine that this system can be operated over a wavelength range of about 170nm in most regions without severely increasing the image size. Thus the ADORE system is well suited for VCSELs arrays spanning a range of wavelengths and potentially LED sources.




Figure 73. Theoretical image size dependence on wavelength and focal lengths for 345 μm focal length microlenses operated at 850nm. Range is from unit design image size (white) to >2.5 times image size (black).

The chromatic dependence of the ADORE system was measured for an individual optical channel. The diffractive elements were fabricated courtesy of the Honeywell/ARPA CO-OP diffractive optics foundry run. Each piece designed by the authors provides an experimental test kit of several sets of lenses and lens arrays. The binary phase level microlenses were chosen to have focal lengths of 345 μm when operated at 850nm and have 125 μm apertures. The relay lens chosen was an eight phase level lens with a focal length of 16mm and a 2mm aperture. Micropositioners were used to align the setup. To demonstrate the diminished chromatic sensitivity of the ADORE system, a single mode fiber coupled to a tunable Ti:Sapphire laser was positioned at the location of the proposed VCSEL source. The tunable range of the laser was limited to 770-870nm.

Figure 74 shows the beam size at the second microlens as measured by a knife-edge beam scanner. Typical ADORE behavior shows the beam size decreasing slightly on either side of the design




Figure 74. Measured beam size at 2nd micro-lens as a function of wavelength compared to simulated values.




Figure 74. Image size as a function of wavelength compared to simulated values.

wavelength and then quickly increasing beyond about $\delta\lambda/\lambda = \pm 0.12$. This is desirable since the beam must be confined within the lens aperture to avoid crosstalk. Figure 74 shows that the magnified image size is only modestly influenced by wavelength.

In a further experiment, an ADORE setup was used to successfully create a connection between two linear fiber arrays of a VCSEL-driven parallel optical data link. A follow-up experiment is planned that removes the fiber and creates a direct connection between the linear VCSEL and receiver arrays.

One scheme for using an ADORE module on a conventional electronic circuit board is shown in figure 75. The microlenses are attached with micron-level accuracy to the integrated VCSEL/CMOS circuit. This accuracy will be achieved using bonding techniques similar to those developed for VCSEL-CMOS attachment. The chip is then mounted on the circuit board. The relay lens component, constructed as a solid unit, would then be added using either passive or active alignment depending on requirements. The expected alignment accuracy should be of the order of a few to tens of microns. Turning mirrors allow the connections to be made across a horizontal surface. If the VCSELs/receivers were arranged on 125 μm spacings, a total of 1024 optical channels could be supported in a 4mm x 4mm cross sectional area. In addition, planar integration architectures[168] can also use a similar technique to provide chip-to-chip interconnection.

In summary, high throughput optoelectronic circuits will require strategic integration of the optomechanical infrastructure with standard electronic packaging schemes. The ADORE interconnection concept presents an opportunity to create inexpensive circuit board level optical interconnections that are relatively insensitive to VCSEL spectral performance and require relaxed alignment tolerances.




Figure 75. Chip to chip ADORE interconnection after the chips are attached to the circuit board. An alternative approach would be to incorporate the ADORE module into the circuit board itself.

5. RECOMMENDATIONS

The technology has advanced to the point where products may be placed in the field without major new inventions. The progress made under this contract, particularly in terms of evolution in the optoelectronic VLSI technology, but also in the optomechanics and test instrumentation, certainly moved the state of the art forward toward that goal. The improvements outlined in system 7, particularly the need for small size and direct interfacing to electronic circuit boards, was also a step in the right direction. Beyond system 7, several technological issues remain, notably the ability to work with standard parts, be it fibre-channel, ethernet or other standards.

It is unlikely commercialization of the technology will take place without government sponsorship. We have seen a change in emphasis at Lucent from high performance systems to low cost systems, with the most important constraint being "time to market." Indeed smaller firms have been driven by this for years, large firms either will be or will not survive. We have analyzed many systems in telecommunications and data networking; we have yet to find one that *requires* free space optical interconnects and optoelectronic VLSI, yet all of these systems, particularly as their capacities grow, would benefit greatly from these technologies if it were available to put into their products. The benefits come mainly in cost, space, and power dissipation, metrics that are much more important in military avionics than on the floor of a central office switch.

If we could recommend one thing to DARPA and the other government sponsors it would be to continue to fund *system level* research in optical interconnects. Within the past several years, that research has declined, with much emphasis being placed on surface emitting lasers particularly at the device level. It may still be several years, indeed if at all, before a VCSEL OE-VLSI platform exists with performance, yields and reliability comparable to the modular based platform that we (and others) have developed. But to a major degree, the choice of output device from the switching array is so minor compared to the other issues, particularly receivers, fiber arrays, optics, mechanics, and reliability. It is our hope that research will continue to improve these aspects, so that free-space optically interconnected systems may someday be a reality.

6. REFERENCES

1. <http://cell-relay.indiana.edu>
2. D. A. B. Miller, "Physical reasons for optical interconnections," *Intl. J. Optoelectronics* 11 (3), 155-168 (1997).
3. W.J. Dolly, and S. Poulton, "Transmitter Equalization for 4Gbps Signalling," *IEEE Micro* 17, 48-56 (1997).
4. D. A. B. Miller, "Physical and Systems Motivations for Smart Pixels", *IEEE LEOS Topical Meeting on Smart Pixels*, Keystone Colorado, August 1996
5. M.E. Prise, N.C. Craft, R.E. LaMarche, M.M. Downs, L.A. D'Asaro, and L.M.F. Chirovsky, "Cascaded operation of arrays of symmetric self-electrooptic-effect devices," *Appl. Optics*, 30, 2841-2843 (1991).
6. J.W. Goodman, "A short history of the field of optical computing," in *Optical Computing, Proceedings of the 34th Scottish Universities Summer School in Physics, 1988*, B.S. Wherrett and F.A.P. Tooley, eds., IOP Publishing Ltd. (Adam Hilger) Bristol England and New York, NY.
7. E. Kerbis, T.J. Cloonan, and F.B. McCormick, "An all optical realization of a 2x1 free-space switching node," *IEEE Photon. Tech. Lett* 2, 600-602 (1990).
8. G.W. Richards and F. K. Hwang, "Extended Generalized Shuffle Networks: Sufficient conditions for strictly non-blocking operation," *Networks* (to be published)
9. J. Jahns and M. Murdocca, "Crossover networks and their optical implementation," *Appl. Opt.* 27, 3155 (1988).
10. M.E. Prise, M.M. Downs, F.B. McCormick, S.J. Walker and N. Streibl, "Design of an optical digital computer," in *Optical Bistability IV*, (W. Firth, N. Peyghanbarian, and A. Tallet eds.) Les Editias De Physique, Les Ulis Cedex France, (1988).
11. F.A.P. Tooley and S. Wakelin, "Design of a symmetric self-electrooptic effect device cellular image processor," *Appl. Opt.* 32, 1850 (1993).
12. T. J. Cloonan, M. J. Herron, F. A. P. Tooley, G. W. Richards, F. B. McCormick, E. Kerbis, J. L. Brubaker, and A. L. Lentine, "An all optical implementation of a 3D Crossover network," *IEEE Photon. Technol. Lett.*, 2, 438-440 (1990)
13. A. L. Lentine, H. S. Hinton, D. A. B. Miller, J. E. Henry, J. E. Cunningham, and L. M. F. Chirovsky, "Symmetric self-electro-optic effect device: Optical set-reset latch, differential logic gate, and differential modulator/detector," *IEEE J. Quantum Electron.*, QE-25, 1928-1936 (1989)
14. A. L. Lentine, D. J. Reiley, R. A. Novotny, R. L. Morrison, J. M. Sasian, M. G. Beckman, D. B. Buchholz, S. J. Hinterlong, T. J. Cloonan, G. W. Richards, and F. B. McCormick, "ATM distribution network using an optoelectronic VLSI switching chip," *Appl. Opt.* 36 1804-1814 (1997)
15. A. L. Lentine, F. A. P. Tooley, S. L. Walker, F. B. McCormick, R. L. Morrison, L. M. F. Chirovsky, M. W. Focht, J. M. Freund, G. D. Guth, R. E. Leibenguth, G. J. Przybylek, and L. E. Smith, L. A. D'Asaro and D. A. B. Miller, "Logic self electro-optics effect devices: Optoelectronic multiport logic gates, multiplexers, demultiplexers and shift registers," *IEEE J. Quantum. Electron.*, 28, 1539-1553 (1992)
16. A. L. Lentine, R. A. Novotny, T. J. Cloonan, L. M. F. Chirovsky, L. A. D'Asaro, S. P. Hui, M. W. Focht, J. M. Freund, G. D. Guth, R. E. Leibenguth, K. G. Glogovsky, and T. K. Woodward, "4 x 4 Arrays of FET-SEED embedded control 2 x 1 optoelectronic switching nodes with electrical fan-out," *IEEE Photon. Technol. Lett.*, 6, 1126-1129 (1994).
17. for a summary see F. A. P. Tooley, "Challenges in optically interconnecting electronics," *IEEE JSTQE* 2 3-13 (1996)

18. M. J. Goodwin and A. J. Mosely, "The application of optoelectronic technologies to high performance electronic processor interconnects," *Opt. Quantum E.* 9 S455-S470 (1994)
19. K. W. Goossen, J. A. Walker, L. A. D'Asaro, S. P. Hui, B. J. Tseng, R. E. Leibenguth, D. Kossives, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller, "GaAs MQW modulators integrated with silicon CMOS," *IEEE Photon. Technol. Lett.*, 7, 360-362 (1995)
20. see <http://www.bell-labs.com/projects/oevlsi>
21. T. K. Woodward, A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Y. Jan, L. A. D'Asaro, L. M. F. Chirovsky, S. P. Hui, B. J. Tseng, D. P. Kossives, D. W. Dahringer, and R. E. Leibenguth, "1 Gb/s two beam transimpedance smart pixel optical receiver made from hybrid GaAs MQW-modulators bonded to 0.8 mm silicon CMOS," *IEEE PTL*. 8 442-444 (1996).
22. T. L. Worchesky, R.J. Ritter, R. Martin and B. Lane, "Large arrays of spatial light modulators hybridized to silicon integrated circuits," *Appl. Opt.*, 35, 1180-1186 (1996)
23. A. L. Lentine, K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Y. Jan, T. K. Woodward, A. V. Krishnamoorthy, S. P. Hui, B. J. Tseng, R. E. Leibenguth, L. M. F. Chirovsky, R. A. Novotny, D. B. Buchholz and R. L. Morrison. "Optoelectronic VLSI switching chip with greater than 1 terabit per second potential I/O bandwidth," *Electronics Lett.* 33 894-895 (1997)
24. Y. Ohiso, K. Tateno, Y. Kohama, A. Wakatsuki, "Flip-chip Bonded 0.85- μ m Bottom-emitting Vertical-cavity Laser Array On An AlGaAs Substrate," *IEEE Photon. Technol. Lett.*, 8 1115-1117 (1986).
25. D. Paananen, J. Wasserbauer, J. Scott, P. Brusenbach, D. Lewis, and G. Simonis, "A 16 x 16 vertical cavity surface emitting laser array module for optical processing and interconnect applications," *OSA Topical meeting on Optics in Computing*, Postdeadline paper PD1 (1997).
26. C. Fan, D. A. Vanbierkom, W. L. Hendrick, and S. C. Esener, "Free-space Optical Interconnection - A Technology Comparison Of Vertical-cavity Surface-emitting Lasers And Multiple-quantum-well Modulators," *Rev.* 3 6A 394 (1996).
27. C. Fan, B. Mansoorian, D. A. Vanbierkom, M. W. Hansen, and S. C. Esener, "Digital Free-space Optical Interconnections - A Comparison Of Transmitter Technologies," *Appl. Opt.* 34 3103-3117 (1995).
28. T. Nakahara, S. Matsuo, S. Fukushima, and T. Kurokawa, "Performance Comparison Between Multiple-quantum-well Modulator-based And Vertical-cavity-surface-emitting Laser-based Smart Pixels," *Appl. Opt.* 35 860-871 (1996)
29. K. W. Goossen, J. E. Cunningham, W. Y. Jan and R. E. Leibenguth, "On the operational and manufacturing tolerances of GaAs/AlGaAs MQW Modulators, *IEEE LEOS Annual Meeting*, 196-197 (1997).
30. R. A. Novotny, "Analysis of smart pixel digital logic and optical interconnections," Ph.D. Thesis, Heriot-Watt Univ. 1997
31. A. L. Lentine, D. A. B. Miller, L. M. F. Chirovsky, and L. A. D'Asaro, "Optimization of absorption in symmetric self electro-optic effect devices: A systems perspective," *IEEE J. Quantum Electron.*, QE-27, 2431-2439
32. K. W. Goossen, J. E. Cunningham and W. Y. Jan, "Stacked Diode Electroabsorption Modulator," *IEEE Photon. Technol. Lett.*, 6 936-938 (1994).
33. T. K. Woodward, A. V. Krishnamoorthy, K. W. Goossen, J. A. Walker, B. Tseng, J. Lothian, S. P. Hui, R. E. Leibenguth, "Modulator-driver circuits for optoelectronic VLSI", *IEEE Photon. Technol. Lett.*, 9 839-841, (1997).
34. A. L. Lentine, L. M. F. Chirovsky and T. K. Woodward, "Optical energy considerations for diode-clamped smart pixel receivers," *IEEE J. Quantum Electron.*, 30, 1167-1171 (1994)
35. L. Boivin, M. C. Nuss, J. Shah, and D. A. B. Miller, "Receiver Sensitivity Improvement By Impulsive Coding," *IEEE Photon. Technol. Lett.*, 9, 684-686 (1997).

36. T. H. Wood, E. C. Carr, B. L. Casper, R. A. Linke, C. A. Burrus, and K. L. Walker, "Bidirectional fibre-optical transmission using a multiple quantum well (MQW) modulator/detector," *Electron. Lett.*, 22 528 (1986)
37. A. V. Krishnamoorthy, T. K. Woodward, K. W. Goossen, J. A. Walker, S. P. Hui, B. T. Tseng, J. E. Cunningham, W. Y. Jan, F. Kiamilev, D. A. B. Miller, "A dual-function detector modulator smart pixel module", *Appl. Optics*, 36 4866-4870, (1997).
38. J. Trezza and J. S. Harris, "2-state Electrically Controllable Phase Diffraction Grating Using Arrays Of Vertical-cavity Phase Flip Modulators, *IEEE Photon. Technol. Lett.*, 8 1211-1213 (1996).
39. L. Wilkinson, S. M. Prince, M. P. Y. Desmulliez, C. R. Stanley, "Fabrication and testing of multiple quantum well binary phase modulator, *OSA Topical Meeting on Spatial Light Modulators* 114-116(1997)
40. H. Yamazaki, M. Yamaguchi and K. Hirabayashi, "Estimation Of The Possible Scale For Holographic Switches With Liquid-crystal Displays ; *Appl. Opt.* 34 1333-1340 (1995).
41. A. Hussain, "DARPA's new vision in optical interconnects," *IEEE High Speed Interconnects Workshop* (1997)
42. J. A. Walker, J. E. Ford, and N. Basavanahally, "Performance and Packaging Implications of a MEMS Based Optical Modulator for WDM Fiber-To-The Home Systems 1997 Proceedings. 47th Electronic Components and Technology Conference (Cat. No.97CH36048), 1997.
43. T. Mukaiharu, N. Ohnoki, Y. Hayashi, N. Hatori, "Excess Intensity Noise Originated From Polarization Fluctuation in Vertical-cavity Surface-emitting Lasers," *IEEE Photon. Technol. Lett.*, 7 1113-1115 (1995).
44. K. Choquette, R. E. Leibenguth, "Control of Vertical Cavity Laser Polarization with Anisotropic Transverse Cavity Geometries," *IEEE Photon. Technol. Lett.*, 6, 40-42 (1994).
45. T. Yoshikawa, H. Kosaka, K. Kurihara, and M. Kajita, "Complete Polarization Control Of 8x8 Vertical-cavity Surface-emitting Laser Matrix Arrays," *Appl. Phys. Lett.* 66 908-910 (1995).
46. D. Vakhshoori, "Symmetry Considerations in Vertical Cavity Surface Emitting Lasers: Prediction of Removal of Polarization Isotropy on (001) Substrates," *Appl. Phys. Lett.* 65 259-261 (1994).
47. T. Mukaiharu, N. Ohnoki Y. Hayashi, N. Hatori, "Polarization Control Of Vertical-cavity Surface-emitting Lasers Using A Birefringent Metal/dielectric Polarizer Loaded On Top Distributed-bragg-reflector," *IEEE J. Sel. Top. in Quantum Electron.* 1 667 (1995)
48. S. J. Schablitsky, L. Zhuang, R. C. Shi and S. Y. Chou, "Controlling Polarization Of Vertical-cavity Surface-emitting Lasers Using Amorphous-silicon Subwavelength Transmission Gratings," *Appl. Phys. Lett.* 69 7-9 (1996)
49. T. A. Strand, B. J. Thibeault, and L. A. Coldren, "Reduced Lateral Carrier Diffusion For Improved Miniature Semiconductor-lasers," *J. Appl. Phys.* 81 3377 (1997)
50. N. K. Dutta, "Power penalty due to timing jitter for lasers modulated without prebias," *Proc. SPIE Vol. 2683*, p. 68-72, *Fabrication, Testing, and Reliability of Semiconductor Lasers*, Mahmoud Fallahi; S. C. Wang; Eds. SPIE vol. 2683 68-72 (1996)
51. P. Schnitzer, U. Fiedler, G. Reiner, B. Weigl, "Bias-free 1-gb/s Data-transmission Using Top Vertical-cavity Surface-emitting Laser-diodes," *IEEE Photon. Technol. Lett.*, 9 693 (1997)
52. W. Nakwaski and M. Osinski, "Available Output Of 2-dimensional Surface-emitting Laser Arrays," *Optical and Quantum Electronics* 29 639 (1997).
53. Y. C. Lee, S. E. Swirhun, W. S. Fu, and T. A. Keyser, "Thermal Management Of VCSEL-based Optoelectronic Modules," *IEEE Trans. on Components, Packaging and Manufacturing Tech. Part B-Advanced Packaging*, 19, 540-547 (1996)
54. J. W. Scott, D. B. Young, B. J. Thibeault, M. G. Peters, and L. A. Coldren, "Design Of Index-guided Vertical-cavity Lasers For Low Temperature-sensitivity, Sub-milliamp Thresholds, And Single-mode Operation *IEEE J.*

Selected Topics in Quantum Electron., 1, 638-648 (1995)

55. Wipiejewski, D. B. Young, B. J. Thibeault, L. A. Coldren, "Thermal Crosstalk In 4x4 Vertical-cavity Surface-emitting Laser Arrays," IEEE Photon. Technol. Lett., 8 980-982 (1996)
56. R. A. Morgan, M. K. Hibbs-Brenner, T. M. Marta, R. A. Walterson, "200-degrees-c, 96-nm Wavelength Range, Continuous-wave Lasing From Unbonded GaAs Mowpe-grown Vertical-cavity Surface-emitting Lasers IEEE Photon. Technol. Lett., 7 441-443 (1995)
57. T. E. Van Eck, P. Chu, W. S. C. Chang, and H. H. Wieder, "Electroabsorption in an InGaAs/GaAs strained-layer multiple quantum well structure," Appl Phys. Lett., vol. 49, 135, 1986.
58. W. Dobbdaere, D. Huang, S. Kalem, and H. Morkoc, Electron. Lett., vol. 24, p. 1239, 1988.
59. T. K Woodward, T. Sizer, D. L Sivco, and A. Y. Cho, "InGaAs/GaAs multiple quantum well modulators for the 1.02-1.07 μm range," Appl Phys. Lett., vol. 57, p. 548, 1990.
60. T. Sizer T. K Woodward, T.-H. Chiu, D. L. Sivco, and A. Y. Cho, "Optoelectronic device mapping using differential imaging techniques," J. Appl Phys., vol. 70, p. 3837, 1991.
61. B. Pezeshki, S. M. Lord, and J. S. Harris, Jr., "Electroabsorptive modulators in GaAs/AlGaAs," Appl Phys. Lett., vol. 59, p. 888, 1991.
62. J. E. Cunningham, K W. Goossen, M. WiDiarns, and W. Y. Jan, "Pseudomorphic InGaAs-GaAsP quantum well modulators on GaAs," Appl. Phys. Lett., vol. 60, p. 727, 1992.
63. K W. Goossen, J. E. Cunningham, and W. Y. Jan, "Defect-free modulator at 1064 nm using a strain-balanced multiquantum well," Electron. Lett., vol. 28, p. 1833, 1992.
64. K W. Goossen, J. E. Cunningham, M. B. Santos, and W. Y. Jan, "Measurement of modulation saturation intensity in strain-balanced, undefected InGaAs/GaAsP modulators operating at 1064nm," Appl. Phys. Lett., vol. 63, p. 515, 1993.
65. K W. Goossen, J. E. Cunningham, and W. Y. Jan, "Exciton electroabsorption in extremely shallow quantum wells," Appl. Phys. Lett., vol. 57, p. 2582, 1990.
66. D. R P. Guy, D. D. Besgrove, L L Taylor, N. Apsloy, and S. J. Bass, "Experimental study of InGaAs-InP MQW electro-absorption modulators," IEEE~ Proceedings, vol. 136, p. 46, 1989.
67. J. Peldman, K W. Goossen, D. A. B. Miller, A. M. Fox, J. E. Cunningham, and W. Y. Jan, "Fast escape of photo-created carriers out of shallow quantum wells," Appl. Phys. Lett., vol 59, p. 66, 1991.
68. W. H. Knox, D. A. B. MiDor, T. C. Damen, D. S. Chemla, A. C. Gossard, and C V. Shank, "Sub-picosecond excitonic electroabsorption in room-temperature quantum wells," Appl. Phys. Lett., vol. 48, p. 864, 1986.
69. D.A.B. Miller, M.D. Feuer, Y.T. Chang, A.S.C. Shunk, J.E. Henry, D.J. Burrow, D.S. Burrows, D.S. Chemla, "Field-Effect Transistor Self-Electrooptic Effect Device: Integrated Photodiode, Quantum Well Modulator and Transistor," IEEE PTL, V1, N3, March 1989.
70. L. A. D'Asaro, L. M. F. Chirovsky, E. J. Laskowski, S. S. Pei, T. K. Woodward, A. L. Lentine, R. E. Leibenguth, M. W. Focht, J. M. Freund, G. D. Guth, and L. E. Smith, "Batch fabrication and operation of GaAs-AlGaAs field effect transistor self electro-optic effect device (FET-SEED) smart pixel arrays," IEEE J. Quantum Electron., 29 670-677 (1993).
71. R.S. Sussmann, R.M. Ash, A.J. Moseley, R.C. Goodfellow, "Ultra- low Capacitance Flip-chip-bonded GaInAs PIN Photodetector for Long Wavelength High-data-rate Fiber Optic Systems," Elec Lett., V21, N14, July 1985]
72. K.W. Goossen, J.E. Cunningham, W.Y. Jan, "GaAs 850-nm Modulators Solder Bonded to Silicon," IEEE PTL V5, 7, 776-778, July 1993.

73. A.L. Lentine, K.W. Goossen, J.A. Walker, L.M.F. Chirovsky, L.A. D'Asaro, S.P. Hui, B.J. Tseng, R.E. Leibenguth, J.E. Cunningham, W.Y. Jan, J.M. Kuo, D.W. Dahringer, D.P. Kossives, D.D. Bacon, G. Livescu, R.L. Morrison, R.A. Novotny, and D.B. Buchholz. "High speed optoelectronic VLSI switching chip with greater than 4000 optical I/O based on flip chip bonding of MQW modulators and detectors to silicon CMOS," IEEE Journal on Selected Topics in Quantum Electronics, 2, 77-84 (1996).
74. R.A. Novotny, A.L. Lentine, D.B. Buchholz, A.V. Krishnamoorthy, "Analysis of Parasitic Front-end Capacitance and Thermal Resistance in Hybrid Flip-chip-bonded GaAs SEED/Si CMOS Receivers," Optical Computing Technical Digest, 207-209, March 1995.
75. A.V. Krishnamoorthy, T.K. Woodward, R.A. Novotny, K.W. Goossen, J.A. Walker, A.L. Lentine, L.A. D'Asaro, S.P. Hui, B.J. Tseng, R.E. Leibenguth, D.P. Kossives, D.W. Dahringer, L.M.F. Chirovsky, G.F. Aplin, F.E. Kiamilev, D.A.B. Miller, "Ring Oscillators with Optical and Electrical Readout Based on Hybrid GaAs MQW Modulators Bonded to 0.8- μ m Silicon VLSI Circuits," Electronics Letters, V31, 22, p1917-1918, 1995.
76. T. Sakurai, K. Tamaru, "Simple Formulas for Two and Three-Dimensional Capacitances," IEEE Trans on Elec. Dev., Vol. ED-30, No2, February 1983.
77. MOSIS wafer acceptance specifications for HP CMOS26B 0.8 micron CMOS bulk wafers.
78. Capacitance values in Table 4.2 were calculated using the MOSIS wafer acceptance specifications for HP CMOS26B 0.8micron CMOS bulk wafers.
79. METAL Two and Three dimensional interconnect simulator, OEA International
80. N.H.E. Weste, K. Eshraghian, "Principals of CMOS VLSI Design," Addison-Wesley, ISBN: 0-201-53376-6, 1993.
81. SPICE is a general purpose simulation program distributed free of charge by ilpsoftware@berkeley.edu
82. A.L. Lentine, F.B. McCormick, R.A. Novotny, L.M.F. Chirovsky, L.A. D'Asaro, R.F. Kopf, J.M. Kuo, G.D. Boyd, "A 2 kbit Array of Symmetric Self-Electrooptic Effect Devices," IEEE PTL, 2, pp51-53, 1990.
83. J.M. Kuo, Y.C. Wang, K.W. Goossen, L.M.F. Chirovsky, S.P. Hui, B.T. Tseng, J. Walker, A.L. Lentine, R.E. Leibenguth, G. Livescu, W.Y. Yan, J.E. Cunningham, L.A. D'Asaro, D.W. Dahringer, D.P. Kossives, D.D. Bacon, R.L. Morrison, R.A. Novotny, D.B. Buchholz, "Large Array of GaAs Modulators and Detectors Flip-Chip Solder Bonded to Silicon CMOS using InGaP as the Selective etch stop layer for GaAs Substrate Removal," Submitted to the Ninth International Conference on Molecular Beam Epitaxy, Aug 5-9, 1996.
84. O. Wada, M. Makiuchi, H. Hamaguchi, T. Kumai, T. Mikawa, "High Performance, High-Reliability InP/GaInAs p-i-n Photodiodes and Flip-Chip Integrated Receivers for Lightwave Communications," IEEE JLT, Vol9, N9, September, 1991.
85. D.A.B. Miller, D.S. Chemla, T.C. Damen, A.C. Gossard, W. Wiegmann, T.H. Wood, C.A. Burrus, "Electric Field Dependence of Optical Absorption near the Bandgap of Quantum Well Structures," Phys. Rev. B, 32, pp.1043-1060, 1985.
86. D.B. Buchholz, A.L. Lentine, R.A. Novotny, "Thermal Considerations in the Design of Opto-Electronic Device Mounts," Photonics in Switching Technical Digest, pp.118-120, March, 1995.
87. T.K. Woodward, W.H. Knox, B. Tell, J.B. Stark, "Experimental Studies of Proton Implanted GaAs AlGaAs Multiple Quantum Well Modulators for Low Photocurrent Applications," IEEE JQE, V30, 12, p2854-2865, 1994.
88. K.W. Goossen, B. Tseng, S.P. Hui, J.A. Walker, R. Leibenguth, L.M.F. Chirovsky, A. Krishnamoorthy, "Multiple Attachment GaAs-On-Si Hybrid Optoelectronic/VLSI Chips," Submitted to IEEE Summer Topicals on Smart Pixels, 1996.
89. R.R. Tummala, E.J. Rymaszewski, "Microelectronics Packaging Handbook," Van Nostrand Reinhold, Ch 4, ISBN:0-442-20578-3, 1989.

90. B. Mrozievicz, M. Bugajski, W. Nakwaski, "Physics of Semiconductor Lasers," North-Holland, ISBN: 0-444-98737-1, 1991.
91. S.L. Chuang, "Physics of Optoelectronic Devices," John Wiley & Sons, ISBN:0-471-10939-8, 1995.
92. The transient thermal effect is under investigation.
93. A.L. Lentine, K.W. Goossen, J.A. Walker, L.M.F. Chirovsky, L.A. D'Asaro, S.P. Hui, B.T. Tseng, R.E. Leibenguth, D.P. Kossives, D.W. Dahringer, D.D. Bacon, T.K. Woodward, D.A.B. Miller, "Arrays of Optoelectronic Switching Nodes Comprised of Flip-Chip-Bonded MQW Modulators and Detectors on Silicon CMOS Circuitry," IEEE Photon. Technol. Lett. Vol 8, 221-223 (1996).
94. J. Davidse, "Analog Electronic Circuit Design," Prentice Hall International, ISBN: 0-13-035346-9, 1991.
95. B. Tell, K.F. Brown-Goebeler, R.E. Leibenguth, F.M. Baez, Y.H. Lee, "Temperature dependence of GaAs-AlGaAs Vertical Cavity Surface Emitting Lasers," APL V60, pp.683-685, February, 1992.
96. R.A. Novotny, E. Kerbis, J.L. Brubaker, N. Basavanhally, J. Freund, "Two-dimensional fiber optical data link using self-electrooptic effect device modulators and OEIC detectors," OSA Annual Meeting, November 1991.
97. P.R. Grey, R.G. Meyer, "Analog Integrated Circuits," second edition, John Wiley & Sons, ISBN: 0-471-87493-0, 1984.
98. A.L. Fisher, N. Linde, "A 50-Mbit/s CMOS Optical Transmitter Integrated Circuit," IEEE JSSC, Vol. SC-21, No. 6, December 1986.
99. SPICE is a general purpose circuit simulation program distributed free of charge by ilpsoftware@berkeley.edu.
100. M. J. Goodwin and A. J. Moseley, "The application of optoelectronic technologies to high performance electronic processor interconnects," Optical and Quantum Electronics 26 S455-S470 (1994)
101. K. Y. Eng, M. J. Karol, and Y. S. Yeh, "A growable packet switch architecture: design principles and applications," IEEE Trans. on Commun., 40, 423-430 (1994)
102. T. J. Cloonan and G. W. Richards, "Terabit per second packet switch having distributed out-of-band control of circuit and packet switching communications," US Patent Number 5537403.
103. D. J. Reiley, M. G. Beckman, and J. M. Sasian, "Optomechanical design for free space optical switching," Optoelectronic Packaging, SPIE Proceedings 2691 (1996)
104. F. B. McCormick, T. J. Cloonan, A. L. Lentine, J. M. Sasian, R. L. Morrison, R. A. Novotny, M. G. Beckman, S. L. Walker, M. J. Wojcik, S. J. Hinterlong, R. J. Crisci, and H. S. Hinton, "A 5 stage free space optical switching network using FET-SEED smart pixel arrays," Appl. Opt.,33, 1601-1618 (1994)
105. T. J. Cloonan, G. W. Richards, and A. L. Lentine, "Optical implementation of a parallel out of band controller for large ATM switch applications," Optical interconnects in broadband switching architectures, SPIE Proceedings vol. 2692 (1996)
106. D. J. Reiley and J. M. Sasian, Applied Optics (to be published)
107. J. M. Sasian, et. al., "Fabrication of Fiber Arrays for Optical Computing and Switching Systems," OSA conference on Optical Computing, Edinburgh, Scotland, 8/94
108. R. L. Morrison, S. G. Johnson, A. L. Lentine, and W. Knox, "Design and Demonstration of a High Speed Multichannel Optical Oscilloscope," Appl. Opt.36, 1187-1194 (1996)
109. A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, T. K. Woodward, J. E. Ford, G. F. Aplin, L. A. D'Asaro, S. P. Hui, B. Tseng, R. E. Leibenguth, D. Kossives, D. Dahringer, L. M. F. Chirovsky, and D. A. B. Miller, "3-D integration of MQW SEED modulators over active sub-micron CMOS circuits: 375 Mb/s transimpedance receiver-transmitter circuit," IEEE Photon. Technol. Lett, 7, 1288-1290 (1995)
110. M. Bazes, "Two novel fully differential complementary self-biased CMOS differential amplifiers," IEEE Jour-

128. Sasian, J. M., "Objective lens for a free-space photonic switching system," U.S. Patent # 5,353,164, 10/4/94
129. J. M. Sasian, F. B. McCormick, R. Webb, R. J. Crisci, K. O. Mersereau, R. P. Stawicki, Design, Assembly, and testing of an objective lens for a free-space photonic switching system," *Optical Engineering* 32 (08), p. 1971-1878, August, 1993
130. Warren J. Smith, "Fundamentals of the Optical Tolerance Budget," International Lens design Conference, George N. Lawrence, Editor, Proc. SPIE 1354, 474-481,(1990).
131. Schott Glass Technologies, Duryea, PA
132. Kreischer Optics, Inc., McHenry, IL
133. R. E. Fischer, "Optimization of lens designer to manufacturer communication," in *International Lens Design Conference*, George N. Lawrence, Editor, Proc. SPIE 1354, pp. 506-522 (1990)
134. W. J. Smith, "Fundamentals of the optical tolerance budget," in *International Lens Design Conference*, George N. Lawrence, Editor, Proc. SPIE 1354, pp. 474-481 (1990)
135. K. P. Thompson, L. Hoyle, "Key to cost effective optical systems: Maximize the number of qualified fabricators," in *OSA Proceedings of the International Optical Design Conference*, G.W. Forbes, ed. (Optical Society of America, Washington, DC 1994), Vol. 22, pp 208-213
136. R. Shannon, OSA Handbook
137. R.L. Morrison, M.J. Wojcik, D.B. Buchholz, "Nonseparable, surface-relief gratings that generate large, arbitrary intensity beam arrays," in *Optoelectronic Interconnects III*, R.T. Chen, H.S. Hinton, Editors, Proc. SPIE 2400, 32-41 (1995).
138. D. Vukobratovich, "Rugged yet lightweight: how can we achieve both in optical instruments?" in *SPIE Proceedings Optomechanical Design*, July 1992
139. E. A. Avallone, T. Baumeister, "Marks Standard Handbook for Mechanical Engineers," McGraw-Hill, 1987
140. J. M. Sasian, et. al., "Fabrication of Fiber Arrays for Optical Computing and Switching Systems," OSA conference on Optical Computing, Edinburgh, Scotland, 8/94
141. P.R. Yoder, "Advanced considerations of the lens-to-mount interface, in *SPIE Proceedings Optomechanical Design*, July 1992
142. R. L. Morrison, et. al., "Beam array generation and holographic interconnections in a free-space optical switching network," *Applied Optics*, vol 32, No. 14, 5/10/93, pp25512-2518.
143. C. M. Miller. "A fiber-optic-cable connector," *The Bell Technical Journal* 54 (9), 1547-1555,1975.
144. U. Dan~er, P. Kipfer. K. Zul1, J. Lindolf, and J. Schwider, "High precision two dimensional fibre-array in silicon V-groove technique,~ *Angewandte Optic*, Physicalisches Institut der Universitat Erlangen, Annual Report 1992.
145. A. Sasald, T. Baba. and K. Iga. "Put-in microconnectors for alignment-free coupling of optical fiber array," *IEEE Photonics Technology Letters*, 4 (5),908-910, 1992.
146. G. A. Koepf and B. J. Markey, "Fabrication and characterization of a 2-D fiber array," *Applied Optics* 23 (2) 3515-3516, 1984.
147. N. Basavanhally, "Opto-mechanical alignment and assembly of 2D-array components,~ *Technical Digest of the IEEE Princeton Second Sarnoff Symposium*, March 26, 1992.
148. G. M. Proudley, C. Stace, White, "Fabrication of 2-D Fibre-Optic Arrays for an Optical Crossbar Switch," Submitted to *Optical Engineering*.
149. K. Koyabu, F. Ohira, T. Yamamoto, and S. Matsuo, "Novel high-density collimator module," in *Technical*

111. A. V. Krishnamoorthy, P. J. Marshand, F. E. Kiemelev, and S. C. Esener, "Grain size considerations for optoelectronic multistage interconnection networks," *Appl. Opt.* 31 5480-5506 (1992)
112. D. A. B. Miller, D. S. Chemla, T. C. Daman, A. C. Gossard, W. Weigmann, T. H. Wood, C. A. Burrus, "Novel hybrid optically bistable switch: The quantum well self-electro-optic effect device," *Appl. Phys. Lett.* 45, pp. 13-15 (1984)
113. D. A. B. Miller, D. S. Chemla, S. Schmitt-Rink, "Electric Field Dependence of Properties of Semiconductor Quantum Wells: Physics and Applications", in *Optical Nonlinearities and Instabilities in Semiconductors*, H. Haug Ed., Academic Press, Inc., 1988
114. S. M. Sze, *Physics of Semi Conductor Devices*, 2nd Ed., John Wiley and Sons, 1981 pp.15
115. R. A. Novotny, "Parallel optical data links using VCSELs," *Proceedings SPIE*, vol. 2147, *Vertical-Cavity Surface-Emitting Laser Arrays*, pp. 140-149 (1994)
116. D.A.B. Miller, M.D.Feuer, Y.T. Chang, A.S.C. Shunk, J.E. Henery, D.J. Burrow, D.S. Barrow, D.S. Chemla, "Field-Effect Transistor Self-Electrooptic Effect Device: Integrated Photodiode, Quantum Well Modulator and Transistor," *IEEE PTL*, V1, N3, March 1989
117. F. B. McCormick, T. J. Cloonan, F. A. P. Tooley, A. L. Lentine, J. M. Sasian, J. L. Brubaker, R. L. Morrison, et. al., "A six-stage digital free-space optical switching network using SEEDs," *Appl. Opt.*, 32, 5153-5171 (1993).
118. F. B. McCormick, F. A. P. Tooley, T. J. Cloonan, J. L. Brubaker, A. L. Lentine, R. L. Morrison, S. J. Hinterlong, M. J. Herron, S. L. Walker, and J. M. Sasian, "Experimental investigation of a free-space optical switching network using S-SEEDs," *Appl. Opt.*, 31, 5431-5446 (1992)
119. Guilfoyle, Peter S.; Zeise, Frederick F.; Stone, Richard V., "Diffractive optical interconnect element design for digital switching," *Proc. SPIE Vol. 2152*, p. 164-172, *Diffractive and Holographic Optics Technology*, Ivan Cindrich; Sing H. Lee; Eds.06/1994
120. Tooley, Frank A. P.; Wakelin, S.; Taghizadeh, Mohammad R., "Interconnects for a symmetric-self-electro-optic-effect-device cellular-logic image processor," *Appl. Opt.*, vol. 33, issue 8, page 1398 (1994)
121. C.P. Barrett, A.C. Walker, D.T. Neilson, and B.R. Robertson, "Optical Design for Crossbar Systems," in *OSA Proceedings of the International Optical Design Conference*, G.W. Forbes, ed. (Optical Society of America, Washington, DC 1994), Vol. 22, pp.227-231
122. Ludwig, David E.; Carson, John C.; Lome, Louis S., "Optoelectronic interconnects for 3D wafer stacks," *Proc. SPIE Vol. CR62*, p. 64-67, *Optoelectronic Interconnects and Packaging*, Ray T. Chen; Peter S. Guilfoyle; Eds. 01/1996
123. Jahns, J., "Tolerant design of planar optical interconnections," *OSA Proceedings of the International Lens Design Conference*, G. W. Forbes, ed. (Optical Society of America, Washington, DC 1994) Vol. 22, pp232-235
124. Liu, Yung S.; Wojnarowski, R. J.; Hennessy, W. A.; Bristow, Julian P.; Liu, Yue; Peczalski, Andrzej; Rowlette, John R.; Plotts, Alan; Stack, Jared D.; Yardley, James T.; Eldada, L.; Osgood, Richard M.; Scarmozzino, Robert; Lee, Sing H.; Ozguz, Volkan H., "Polymer optical interconnect technology (POINT): optoelectronic packaging and interconnect for board and backplane applications," *Proc. SPIE Vol. CR62*, p. 405-414, *Optoelectronic Interconnects and Packaging*, Ray T. Chen; Peter S. Guilfoyle; Eds. 01/1996
125. Cloonan, T. J., Richards, G. W., "Chuted, growable packet switching arrangement," U.S. Patent #5,345,444, 9/6/1994
126. R. L. Morrison, et. al., "Beam array generation and holographic interconnections in a free-space optical switching network," *Applied Optics*, vol 32, No. 14, 5/10/93, pp25512-2518
127. J.M. Sasian, R. A. Novotny, M. G. Beckman, S. L. Walker, M. J. Wojcik, S. J. Hinterlong, "Fabrication of fiber bundle arrays for free-space photonic switching systems," *Optical Engineering*, 33 (09) pp2979-2985, 9/94

Digest, Conference on Optical fiber Communication International Conference on Integrated Optics and Optical Fiber Communication, 1993 Technical Digest Series, Vol. 4 (Optical Society of America, Washington, D.C. 1993), pp. 2-3.

150. F. B. McCormick, A. L. Lentine, R. L. Morrison, J. M. Sasian, T. J. Cloonan, R. A. Novotny, M. G. Beckman, M. J. Wojcik, S. J. Hinterlong, and D. B. Buchholz, "155 Mb/s operation of a FET-SEED free-space switching network," *Photon. Tech. Lett.* **6**, no. 12, 1479-1481 (1994).
151. J. A. Valdmanis and G. Mourou, "Subpicosecond electrooptic sampling: principles and applications," *J. Quant. Elect.*, QE-22, 69-78 (1986).
152. B. H. Kolner and D. M. Bloom, "Electrooptic sampling in GaAs integrated circuits," *J. Quant. Elect.*, QE-22, 79-94 (1986).
153. A. L. Lentine, L. M. F. Chirovsky, L. A. D'Asaro, R. F. Kopf, and J. M. Kuo, "High speed 2x4 array of differential quantum well modulators," *Phot. Tech. Lett.* **2**, no. 7, 477-480(1990).
154. C.P. Barrett, A.C. Walker, D.T. Neilson, and B.R. Robertson, "Optical Design for Crossbar Systems," in *OSA Proceedings of the International Optical Design Conference*, G.W. Forbes, ed. (Optical Society of America, Washington, DC 1994), Vol. 22, pp.227-231
155. E. Delano, First-order design and the $y-\bar{y}$ diagram, *Appl. Opt.* Vol. 2 No. 12, (1963) pp 1251-1256
156. D. Kessler and R. Shack, $y-\bar{y}$ diagram, a powerful optical design method for laser systems, *Appl. Opt.* vol. 31, no. 15,(1992) p2692-2707
157. R. V. Schack, "Analytic system design with a pencil and ruler - the advantages of the $y-\bar{y}$ diagram," in *Applications of Geometrical Optics II*, W. J. Smith, ed., *Proc. Soc. Photo-Opt. Instrum. Eng.*, **39**, 127-140 (1973)
158. F. J. Lopez-Lopez, "The application of the Delano $y-\bar{y}$ diagram to the optical design," Ph.D. dissertation, University of Arizona, Tucson, AZ 1973)
159. O'Shea, D. C., Elements of Modern Optical Design, John Wiley & Sons, New York (1985)
160. T.H. Jamieson, "Hyperhemispherical Field of View Lens Systems", in *Proceedings of the International Optical Design Conference*, G.W. Forbes, ed. (Optical Society of America, Washington, DC 1994), Vol. 22, pp.398-402
161. T. Mori, U.S. Patent #5,296,971, 1994
162. H. Maeda, "An $f-\theta$ lens system," U.S. Patent #4,436,383, 3/13/1984
163. Shafer, David, "Design Challenges for the 1990's." *International Lens Design Conference*, George N. Lawrence, Editor, *Proc. SPIE* 1354, pp. 608-616 (1990)
164. P. R. Yoder, Mounting Lenses in Optical Instruments, SPIE Optical Engineering Press, Volume TT21, Bellingham, WA 1995
165. For example, F.B. McCormick, et. al., "Optical interconnections using microlens arrays," *Opt. and Quan. Elect.* **24**, S465-477 (1992).
166. S. Sinzinger and J. Jahns, "Variations of the hybrid imaging concept for optical computing applications," *OSA topical meeting on Optical Computing*, (OSA Tech. Digest, Salt Lake City,1995) vol. 10, 183-185.
167. R. Kostuk, J. Yeh, M. Fink, "Distributed optical data bus for boardlevel interconnects," *Appl. Opt.* **32**, 5010-5021 (1993).
168. S. Sinzinger and J. Jahns, "Planar optics for optical interconnections in computers," *Topical meeting on Optical Computing (Tech. Digest, Sendai,1996)*, vol. 1, 20-21.

7. INCLUDED PUBLISHED PAPERS

169. K. W. Goossen, M. B. Santos, J. E. Cunningham and W. Y. Jan, "Independence of absorption coefficient-line-width product for multiple quantum wells with excitons from 850 nm to 1064 nm," *IEEE Photon. Technol. Lett.*, 5 1392-1394 (1993).
170. A. L. Lentine, D. J. Reiley, R. A. Novotny, R. L. Morrison, J. M. Sasian, M. G. Beckman, D. B. Buchholz, S. J. Hinterlong, T. J. Cloonan, G. W. Richards, and F. B. McCormick, "ATM distribution network using an optoelectronic VLSI switching chip," *Appl. Opt.* 36 1804-1814 (1997)
171. K. W. Goossen, J. A. Walker, L. A. D'Asaro, S. P. Hui, B. J. Tseng, R. E. Leibenguth, D. Kossives, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller, "GaAs MQW modulators integrated with silicon CMOS," *IEEE Photon. Technol. Lett.*, 7, 360-362 (1995)
172. A. L. Lentine, K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Y. Jan, T. K. Woodward, A. V. Krishnamoorthy, S. P. Hui, B. J. Tseng, R. E. Leibenguth, L. M. F. Chirovsky, R. A. Novotny, D. B. Buchholz and R. L. Morrison. "Optoelectronic VLSI switching chip with greater than 1 terabit per second potential I/O bandwidth," *Electronics Lett.* 33 894-895 (1997)
173. A.L. Lentine, K.W. Goossen, J.A. Walker, L.M.F. Chirovsky, L.A. D'Asaro, S.P. Hui, B.J. Tseng, R.E. Leibenguth, J.E. Cunningham, W.Y. Jan, J.M. Kuo, D.W. Dahringer, D.P. Kossives, D.D. Bacon, G. Livescu, R.L. Morrison, R.A. Novotny, and D.B. Buchholz. "High speed optoelectronic VLSI switching chip with greater than 4000 optical I/O based on flip chip bonding of MQW modulators and detectors to silicon CMOS," *IEEE Journal on Selected Topics in Quantum Electronics*, 2, 77-84 (1996).
174. R.A. Novotny, A.L. Lentine, D.B. Buchholz, A.V. Krishnamoorthy, "Analysis of Parasitic Front-end Capacitance and Thermal Resistance in Hybrid Flip-chip-bonded GaAs SEED/Si CMOS Receivers," *Optical Computing Technical Digest*, 207-209, March 1995.
175. D.B. Buchholz, A.L. Lentine, R.A. Novotny, "Thermal Considerations in the Design of Opto-Electronic Device Mounts," *Photonics in Switching Technical Digest*, pp.118-120, March, 1995.
176. A.L. Lentine, K.W. Goossen, J.A. Walker, L.M.F. Chirovsky, L.A. D'Asaro, S.P. Hui, B.T. Tseng, R.E. Leibenguth, D.P. Kossives, D.W. Dahringer, D.D. Bacon, T.K. Woodward, D.A.B. Miller, "Arrays of Optoelectronic Switching Nodes Comprised of Flip-Chip-Bonded MQW Modulators and Detectors on Silicon CMOS Circuitry," *IEEE Photon. Technol. Lett.* Vol 8, 221-223 (1996).
177. R. L. Morrison, S. G. Johnson, A. L. Lentine, and W. Knox, "Design and Demonstration of a High Speed Multichannel Optical Oscilloscope," *Appl. Opt.* 36, 1187-1194 (1996)
178. T. J. Cloonan and G. W. Richards, "Terabit per second packet switch having distributed out-of-band control of circuit and packet switching communications," *US Patent Number 5537403*.

Independence of Absorption Coefficient-Linewidth Product to Material System for Multiple Quantum Wells with Excitons from 850 nm to 1064 nm

Keith W. Goossen, *Member, IEEE*, M. B. Santos, J. E. Cunningham, and W. Y. Jan

Abstract—We have measured the absorption coefficient (α) and linewidth (Δ) of the excitons of GaAs/AlGaAs and strain-balanced InGaAs/GaAsP multiple quantum-well modulators with wavelengths from 850 to 1064 nm. We find that α decreases and Δ increases as wavelength increases, but their product, and thus the integrated absorption coefficient, remains roughly constant. Thus, the reduced performance observed for longer wavelength modulators is due to exciton broadening.

THE bandwidth and density limitations of electrical interconnects within a computer may be alleviated by using light beams to communicate information between chips. Several information-processing systems are under development which utilize multiple-quantum well (MQW) p-i(MQW)-n modulators as their primary component. In these systems, information is transmitted off an integrated-circuit chip by imprinting it upon the beam reflected off the modulator. This capability should radically increase the aggregate information flow, especially if surface-normal modulators are used (i.e., the light beams are normal to the chip surface), since then two-dimensional arrays with thousands of elements may be formed.

Because of its ability to provide high power with high spectral and spatial quality, the Nd:YAG laser has been considered as a possible light source for these systems. Since the GaAs substrate is a better candidate for large-scale integration than InP, much attention has been focused on InGaAs/GaAs MQW's for this application [1]–[5]. However, since the MQW must be at least 1 μm -thick for useful surface-normal modulation, strain-relief is bound to occur in this system. This relaxation of the lattice results in dislocations which propagate upward resulting in a striated surface [4]. This surface roughness results in unwanted diffraction of the light beams. In addition, the defects make the integration of GaAs transistors problematic.

We have shown that the InGaAs/GaAsP material system may be used to grow undefected MQW's on GaAs for modulators at 1.064 μm [6], [7]. This is because the addition of phosphorous to the barrier results in negative

strain in the barrier balancing the positive strain in the InGaAs well, allowing in principle any number of periods to be grown without having any net strain buildup. Our devices are optically smooth and the sharpness of X-ray-scattering spectra indicate that no relaxation of the lattice occurs. This material system can provide MQW's with band gaps anywhere from 870 nm to 1.064 μm . We have also shown that modulation-saturation is high for modulators using the MQW's, allowing operation as high as tens of kilowatts per square centimeter [8].

It has typically been found by ourselves and others that the absorption coefficient of longer wavelength MQW's is reduced, by about a factor of two for excitons at 1064 nm, compared to GaAs/AlGaAs MQW's with excitons at 850 nm. Since modulators rely on changes in absorption coefficient, the corresponding modulation is also reduced by about a factor of two. Before this work, it was not clear whether the reduction in absorption coefficient was intrinsic to the long-wavelength material systems or due to broadening of the excitons. Here we show that the latter is the case. We measured absorption coefficient (α) and linewidth (Δ) of samples with excitons from 850 to 1064 nm, and found the product of α (when normalized to the total well thickness) and Δ to be roughly constant. Thus, the integrated absorption is constant, and by reducing Δ of long-wavelength modulators their performance should become equal with that of GaAs/AlGaAs devices. Of course, this work does not show whether the broadening is an intrinsic effect (e.g., due to alloy scattering, or possibly, as discussed below, ultrashort exciton lifetimes).

Five long-wavelength p-i(MQW)-n modulators were grown using the InGaAs/GaAsP system on n-type GaAs substrates using Gas-source molecular-beam epitaxy. The mole fraction of In was varied from 0.11 to 0.24 to produce excitons ranging from 920 nm to 1064 nm. Correspondingly, the P content in the barrier was adjusted from 0.6 to 0.75 to maintain a strain balanced condition, resulting in defect free samples, which was checked by X-ray diffraction. The substrates were not rotated during growth, resulting in each sample covering a range of wavelengths. In Fig. 1, the samples from the same substrate are identified by an ellipse about their spectra. Each sample had 50 wells which were 95 Å wide for the samples with excitons

Manuscript received August 24, 1993; revised October 4, 1993.
The authors are with AT & T Bell Laboratories, Holmdel, NJ 07733.
IEEE Log Number 9214204.




Fig. 1. Absorption spectra of various p-i(MQW)-n modulator samples at 0 volts bias (normalized to entire intrinsic width). Spectra grouped by ellipses are from the same (unrotated) wafer.

below 990 nm, 90 Å for the sample with excitons between 1000 and 1040 nm and 85 Å-wide for the 1064 nm sample. The barrier width was 60 Å for all the samples except the 1064 nm sample, which had 65 Å barriers. Atop each sample, a 5000 Å-thick p-type GaAs layer was grown. 200 × 200 μm mesas were etched on each sample, and gold contacts made to the p layer. The backside of each sample was polished and antireflection coatings were applied to both the front and back surfaces.

For comparison, p-i(MQW)-n samples were fabricated using the GaAs/AlGaAs system. A sample with a 1 μm-thick bulk GaAs intrinsic layer was fabricated. A sample with a conventional GaAs/Al_{0.3}Ga_{0.7}As MQW with 86 Å wells and 44 Å barriers (60 periods) was made. Finally, a shallow GaAs/Al_{0.02}Ga_{0.98}As MQW sample, [9] with 100 Å wells and 40 Å barriers (56 periods) was made. All these samples had their substrates removed and were antireflection coated for transmission measurements.

The absorption spectra (normalized to the entire intrinsic width of the sample) of all the samples is shown in Fig. 1. For the long-wavelength samples, the transmission was normalized to unity at a detuning of 120 meV from the exciton peak. For the GaAs/AlGaAs samples, slight residual Fabry-Perot fringes made establishing the absorption baseline somewhat nonquantitative, and can be observed in Fig. 1. As can be seen in Fig. 1, the absorption coefficient is greatly reduced for the long wavelength samples compared to the deep GaAs/AlGaAs MQW. This is partly due to the fact that the barriers are wider for the long wavelength samples, adding optically inert material to the MQW. In Fig. 2, the absorption coefficient, normalized only to the total well thickness in each sample (α) is plotted (open squares). This parameter should more fairly compare samples as if their barrier widths were equal. Note, however, that reducing the barriers of the strain balanced long wavelength samples would require higher strain (more P) in the barriers in order to maintain the strain-balanced condition.

We observe in Fig. 2 that α for the long wavelength




Fig. 2. Plot of absorption coefficient (normalized now to just the total well width, open squares), linewidth (x's) and their product for the samples of Fig. 1. The integrated absorption of the excitons of all the samples is roughly the same.

samples is still greatly reduced compared to the deep GaAs/AlGaAs MQW. However, we also plot in Fig. 2 the linewidths (HWHM, Δ) of the samples' excitons. The linewidth of the long wavelength samples is much larger than that of the deep GaAs/AlGaAs MQW. When the product of absorption coefficient and linewidth ($\alpha\Delta$) is plotted, we see that it remains roughly constant for all the samples. There is a slight rise in the product as wavelength increases, which may be due to increased electron-hole overlap due to the higher barriers, although since the shallow wells show strong excitons with nearly zero confinement, this explanation may be incorrect. Since this product represents the spectrally integrated absorption, this value should remain unchanged if the linewidth of the long wavelength samples were somehow reduced. Therefore we indicate here that if the linewidths of long-wavelength MQW excitons can be reduced to that of GaAs/AlGaAs MQW's, their corresponding absorption coefficients should increase to become comparable also, since the integrated absorption should be constant.

Note that the constancy of the $\alpha\Delta$ product does not extend to other material systems. In [10], $\alpha = 0.94 \text{ 1}/\mu\text{m}$ (normalized to total well width) and Δ (HWHM) = 4.5 meV, giving an $\alpha\Delta$ product of 4.23 meV/μm, less than half of what we find for our samples. (Thus it would appear that InP/InGaAs MQWs have reached the limit of their performance.)

The exciton broadening may be inhomogeneous in nature, i.e., caused by nonuniformities in the sample. The nonuniformities can be in the form of, for example, alloy fluctuations, or as another example, interface roughness. If the latter is the dominant cause, then reduced linewidth may be achieved by improved crystal growth. However, the former is usually thermodynamic in nature, and so probably cannot be controlled greatly by growth conditions. If this is the case, it may not be possible to ever achieve the linewidth of GaAs/AlGaAs MQW's, which have nonallow wells

The exciton broadening may be homogeneous in nature, i.e., caused by ultrashort exciton lifetimes. This is of some concern due to fact that we recently observed large saturation intensities in InGaAs/GaAsP MQW modulators at 1064 nm [8]. This would tend to indicate that fast exciton ionization (escape from the well) is occurring in these samples, an astonishing possibility give the large-band offsets in the system (about 1 eV). However, the effect of lifetime on broadening is probably only small. For shallow GaAs/AlGaAs MQW's, escape from the well has been shown to occur with a single phonon collision [11], which has a time constant of 300 femtoseconds [12]. Since the lifetime induced broadening is given by $\Delta_{\text{life}} = \hbar\tau/(2\tau)$, the linewidth of shallow quantum wells should be at least 1.1 meV broader than that of deep wells GaAs/AlGaAs wells, which have much longer escape times. We see in Fig. 2 that the linewidth of shallow wells is indeed about 2.5 meV broader than deep wells. Since escape from the well probably occurs faster in shallow quantum wells than any other sample, it is unlikely then that more than 1–2 meV of exciton broadening can be attributed to lifetime effects for the long-wavelength samples.

In conclusion, we have shown that the product of absorption coefficient (when normalized to the total well thickness) and exciton linewidth is roughly constant for MQW modulators with exciton wavelengths from 850 to 1064 nm in the GaAs/AlGaAs and InGaAs/GaAsP material systems. Therefore, reducing the linewidth of the long-wavelength samples should result in an increase of their absorption coefficient to that observed for GaAs/AlGaAs. If the broadening is not due to intrinsic effects (e.g., alloy fluctuations or ultrashort exciton lifetimes), improvement of crystal growth could result in

longer wavelength-modulators having performance at least approaching that of 850 nm modulators.

REFERENCES

- [1] T. E. Van Eck, P. Chu, W. S. C. Chang, and H. H. Wieder, "Electroabsorption in an InGaAs/GaAs strained-layer multiple quantum well structure," *Appl. Phys. Lett.*, vol. 49, p. 135, 1986.
- [2] W. Dobbelaere, D. Huang, S. Kalem, and H. Morkoc, *Electron. Lett.*, vol. 24, p. 1239, 1988.
- [3] T. K. Woodward, T. Sizer, D. L. Sivco, and A. Y. Cho, "InGaAs/GaAs multiple quantum well modulators for the 1.02–1.07 μm range," *Appl. Phys. Lett.*, vol. 57, p. 548, 1990.
- [4] T. Sizer, T. K. Woodward, T.-H. Chiu, D. L. Sivco, and A. Y. Cho, "Optoelectronic device mapping using differential imaging techniques," *J. Appl. Phys.*, vol. 70, p. 3837, 1991.
- [5] B. Pezeshki, S. M. Lord, and J. S. Harris, Jr., "Electroabsorptive modulators in InGaAs/AlGaAs," *Appl. Phys. Lett.*, vol. 59, p. 888, 1991.
- [6] J. E. Cunningham, K. W. Goossen, M. Williams, and W. Y. Jan, "Pseudomorphic InGaAs-GaAsP quantum well modulators on GaAs," *Appl. Phys. Lett.*, vol. 60, p. 727, 1992.
- [7] K. W. Goossen, J. E. Cunningham, and W. Y. Jan, "Defect-free modulator at 1.06 μm using a strain-balanced multiquantum well," *Electron. Lett.*, vol. 28, p. 1833, 1992.
- [8] K. W. Goossen, J. E. Cunningham, M. B. Santos, and W. Y. Jan, "Measurement of modulation saturation intensity in strain-balanced, undefected InGaAs/GaAsP modulators operating at 1.064 μm ," *Appl. Phys. Lett.*, vol. 63, p. 515, 1993.
- [9] K. W. Goossen, J. E. Cunningham, and W. Y. Jan, "Exciton electroabsorption in extremely shallow quantum wells," *Appl. Phys. Lett.*, vol. 57, p. 2582, 1990.
- [10] D. R. P. Guy, D. D. Besgrove, L. L. Taylor, N. Apsley, and S. J. Bass, "Experimental study of InGaAs-InP MQW electro-absorption modulators," *IEEE Proceedings*, vol. 136, p. 46, 1989.
- [11] J. Feldman, K. W. Goossen, D. A. B. Miller, A. M. Fox, J. E. Cunningham, and W. Y. Jan, "Fast escape of photocreated carriers out of shallow quantum wells," *Appl. Phys. Lett.*, vol. 59, p. 66, 1991.
- [12] W. H. Knox, D. A. B. Miller, T. C. Damen, D. S. Chemla, A. C. Gossard, and C. V. Shank, "Sub-picosecond excitonic electroabsorption in room-temperature quantum wells," *Appl. Phys. Lett.*, vol. 48, p. 864, 1986.

Asynchronous transfer mode distribution network by use of an optoelectronic VLSI switching chip

A. L. Lentine, D. J. Reiley, R. A. Novotny, R. L. Morrison, J. M. Sasian, M. G. Beckman, D. B. Buchholz, S. J. Hinterlong, T. J. Cloonan, G. W. Richards, and F. B. McCormick

We describe a new optoelectronic switching system demonstration that implements part of the distribution fabric for a large asynchronous transfer mode (ATM) switch. The system uses a single optoelectronic VLSI modulator-based switching chip with more than 4000 optical input-outputs. The optical system images the input fibers from a two-dimensional fiber bundle onto this chip. A new optomechanical design allows the system to be mounted in a standard electronic equipment frame. A large section of the switch was operated as a 208-Mbits/s time-multiplexed space switch, which can serve as part of an ATM switch by use of an appropriate out-of-band controller. A larger section with 896 input light beams and 256 output beams was operated at 160 Mbits/s as a slowly reconfigurable space switch. © 1997 Optical Society of America

1. Introduction

In the past few years, the demand for telecommunications services beyond voice telephony has skyrocketed. For the growth of these services to continue at this rate, cost-effective means of transporting and switching large amounts of information must be found. Although fiber-optic transmission has significantly reduced the cost of transmission, switching high-bandwidth signals remains expensive.

Although totally electronic-based switching systems are certainly feasible for these high-bandwidth systems, fiber-optic connections between frames or racks of equipment have the potential to reduce the cost of these systems. As an example, one can envision fiber-optic data links connecting the line units that receive and transmit data from the outside world with an electronic switching fabric. Optical data links can perform the optical-to-electrical conversions. Multiple optical data links can be electrically

connected with electronic switching chips on a printed circuit board.

As the demand for bandwidth increases, several hundred to several thousand optical fibers might be connected to the switching fabric. Discrete optical data links and parallel data links with up to 32 fibers per data link remain an expensive solution to transporting this information as a result of their per-link cost, physical size, and power dissipation. Power dissipation on the switching chips is high because of the need for electronic drivers for the high-speed electrical interconnections between the switching chips and the data links. Integrating the optical-to-electrical conversions directly onto the switching chips permits building lower-cost and higher-density systems.

In this paper, we describe an experimental optoelectronic switching network based on this lower-cost solution. This demonstration differs in many ways from our earlier system experiments:

1. The system uses a new device technology consisting of GaAs/AlGaAs multiple-quantum-well modulators and detectors flip-chip bonded to silicon VLSI circuitry.^{1,2}

2. The system implements part of a new simplified distribution network for the growable packet architecture.^{3,4} The distribution fabric is implemented with a single chip, contrasting with previous systems consisting of cascaded chips.⁵

The authors are with Bell Laboratories, Lucent Technologies, 2000 North Naperville Road, Naperville, Illinois 60566. J. M. Sasian is with the Optical Sciences Center, University of Arizona, Tucson, Arizona 85721; F. B. McCormick is with the Department of Electrical Engineering, University of California at San Diego, and Call/Recall Corporation, San Diego, California 92121.

Received 5 August 1996; revised manuscript received 13 November 1996.

0003-6935/97/081804-11\$10.00/0

© 1997 Optical Society of America

3. The mechanical design of the system uses a plate-pedestal system⁶ that provides superior robustness compared with the slot-plate systems.⁵ This system is mounted in a standard electronic equipment frame.

4. The system contains a single two-dimensional fiber array providing single-mode fibers for the input signals and read beams and providing multimode fibers for the output beams. The optical system images the inputs from the fiber bundle onto the switching chip, provides optical fan-out of the signals from the fibers to the switching chip, and images the outputs from the chip onto the fiber bundle.

5. The aggregate optical input-output (I/O) bandwidth is greater than in previous systems. In two experiments, the aggregate I/O bandwidths (the sum of the *selected* inputs and outputs multiplied by the bit rate) were 26.6 and 81.9 Gbits/s. The throughput (the product of the number of independent input channels and the bit rate) was 2.9 Gbits/s.

In the remainder of this paper, we give an overview of the demonstration system, starting with the architecture and switching-chip design. We then describe the optical system, input lasers, fiber bundle, and mechanical design. Last we describe the experimental results.

2. Architecture

Asynchronous transfer mode (ATM) is a leading approach to routing high-bandwidth signals for the telecommunications networks of the future. As demands for wide-bandwidth services grows, there may be a need for telecommunications switching networks with aggregate capacities beyond 1 Tbit/s. Although there are many approaches to a large-capacity ATM network, the architecture that we describe here is well suited to this task.

Data routed through ATM networks are sent as fixed-length packets or ATM cells. Each cell contains 53 bytes, 48 for the data itself and five for the routing information. ATM switching networks must have memory for storing ATM cells in the event that two or more cells are destined for the same output at the same time. The size of the memory or memories is determined by the traffic statistics, and sufficient memory must be provided so that the probability of an ATM cell's being dropped should be small, perhaps 10^{-12} . Typically, more than 10,000 cells might need to be stored per output in an ATM switching network. A conceptually simple method of building an ATM switch is to use a single large memory, which may contain many memory chips. The incoming ATM cells are sequentially written into the memory and the outgoing cells are sequentially read from the memory in the appropriate order, depending on their destination. This approach works for ATM switches of modest sizes, but as the number of input and output ports increases, the memory-access time decreases. For example, a 2.5-Gbits/s per channel ATM switch with 256 input and 256 output ports (640-Gbits/s aggregate capacity) would

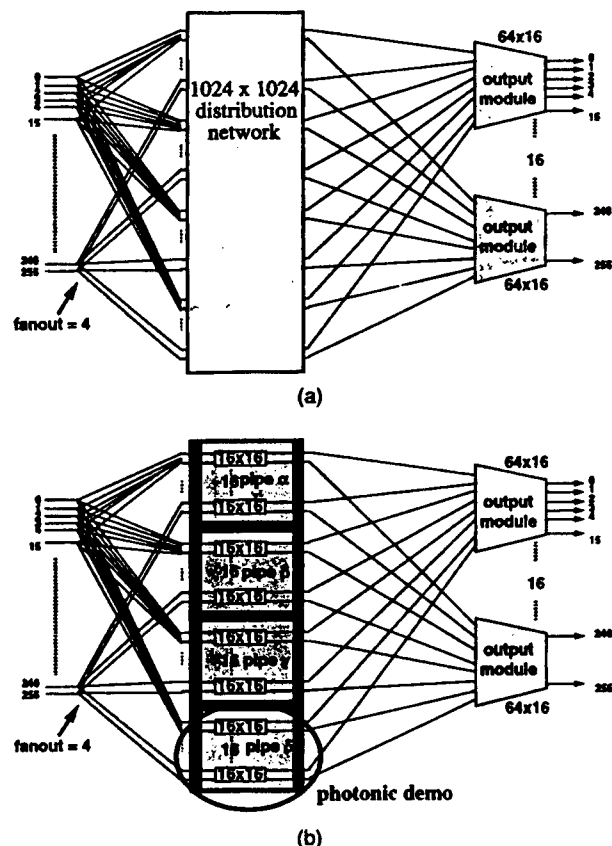


Fig. 1. (a) Growable packet architecture for a 256-input-256-output switching network.³ (b) Distributed interconnected terabit controllable ATM switch architecture implementation.⁴

require subnanosecond access times even if entire ATM cells could be written into and out of the memory in parallel.

The growable packet architecture,³ shown in Fig. 1(a) for a 256×256 ATM network, is one way of partitioning the problem, so that one needs only modest speed memories. The growable packet architecture consists of two subsystems. The first is a distribution network that provides fan-out and distributes the input signals to the second stage, which consists of output-packet modules. The distribution network contains no memory, so it consists of a non-blocking (or low-blocking) interconnection network. The output-packet modules have buffering or memory, as described above, so that ATM cells destined for the same output port can be stored temporarily, and ATM cell loss is minimized. The sizes of the output-packet modules are determined by the cell-loss requirements.

The distribution portion of the network can be a challenge if the network becomes large because of the large number of crosspoints, connections between crosspoints, and calculations that need to be performed to route data through the network during the ATM cell period. A novel architecture, the distributed interconnected terabit controllable ATM switch architecture, has been invented that greatly simplifies




Fig. 2. Switching system environment, consisting of the distribution fabric, input and output interface units, path-hunt processor, output line units, controller interface unit, and output-packet modules.

fies the distribution network.⁴ In Fig. 1(b) we show one implementation of a 256×256 ATM switch using this architecture. The simplified distribution network consists of four groups called pipes, and each pipe consists of sixteen 16×16 switches. This network contains $1/16$ of the number of crosspoints in a 256×1024 crossbar, yet the network has a low blocking probability when combined with the output-packet modules. The reasons for this low blocking probability are that the inputs to the individual switching chips are arranged so that two inputs incident on the same 16×16 switch in one pipe are not incident on the same switch in the other pipes and that the routing algorithm ensures an even distribution of calls through the four pipes.⁷

The distribution network is controlled by an *out-of-band controller*.⁷ In this type of network, the routing information from all the data inputs are routed to a centralized *path-hunt processor* that reads the header information in the ATM cells and calculates the appropriate paths through the distribution network. It must be able to perform these calculations within the ATM cell time (173 ns for 2.5-Gbits/s data and 2.5- μ s for 155-Mbits/s data). The path-hunt processor calculates the paths through the network using global information from all the inputs, performing these calculations during the previous ATM cell period. A major advantage of the architecture is that it allows the calculations of the paths through the network to be calculated in parallel,⁷ so that the path-hunt processor, even for a 1024-input 2.5-Gbits/s per port switching network, can be im-

plemented with logic operating at a clock rate of less than 50 MHz.

The optoelectronic switching demonstration that we describe in this paper implements one pipe of the distribution fabric. Below, we place the demonstration system in the context of other functions that would be implemented in a central office environment, as shown in Fig. 2. The input interface unit contains several functions. First, if the incoming data is from a synchronous optical network (SONET) link, the line unit must provide clock recovery, error detection, SONET pointer processing, and frame delineation. It would also extract the ATM cell, change the routing information contained in the header [virtual path identifier/virtual circuit identifier (VPI/VCI) addresses] to a form that is relevant for the path-hunt processor, perform a translation of the VPI/VCI addresses if necessary, and provide a small amount of buffering of ATM cells so that they can be stored temporarily during path-hunt processing. Our particular implementation also would require the input interface unit to insert a guard band in the data during which the switch reconfigures, encode the data so that long strings of ones or zeros do not occur, add parity-check bits, and add a preamble for synchronization at the receiving end. The 424-bit (53-byte) ATM cell at 155 Mbits/s would be transformed to a 576-bit (72-byte) word at 208 Mbits/s. The input interface unit would contain the lasers that provide the inputs to the optoelectronic distribution network.

The output interface units between the distribu-

tion network and the output-packet modules would detect, decode, and resynchronize the inputs and remove the guard band. It would also perform a second translation of the ATM routing information if necessary. An output line unit, located after the output-packet modules, would convert the ATM cells back into the SONET format. A controller interface unit would provide an interface to a PC or workstation that would allow us to program which physical paths through the distribution fabric corresponded to which VPI/VCI values, as well as perform error checking and other diagnostic functions.


For our demonstration, the implementation of one pipe of the distribution network consists of sixteen 16×16 switches, operating at the SONET standard OC-3c rate of 155 Mbits/s. These 16 switches are realized on a single optoelectronic VLSI switching chip with optical inputs and outputs. A fiber-bundle array with 16 input fibers and 16 output fibers accesses four inputs and four outputs from four of the 16 switches on the chip. A digital word generator, controlled by a PC, supplies time-multiplexed data inputs, formatted as the 72-byte cells that the input interface unit would generate from the incoming SONET inputs. The word generator also supplies the out-of-band control signals in place of the path-hunt processor and reconfigures the switch between cells. The experimental system demonstrates that ATM cells could be sent through the distribution network with the appropriate input and output interface units. Alternatively, the switch can be configured as a space switch, with input data originating from any source, including digitized video, as we discuss in Section 7.

3. Switching Chip

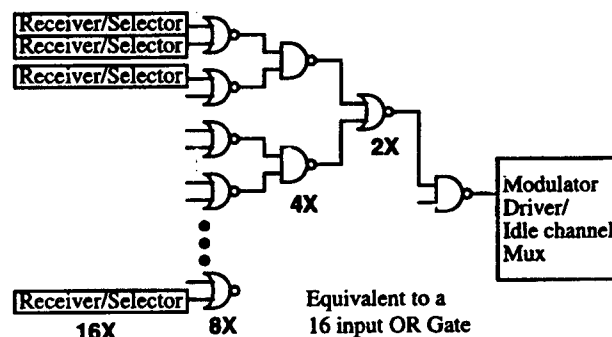
All 16 switches for one pipe of the distribution fabric are implemented on one optoelectronic VLSI chip, as shown in block-diagram form in Fig. 3.⁸ The 16×16 networks are implemented with a passive optical fan-out and electronic fan-in by use of 16×1 multiplexers or switching nodes [Figs. 3(a) and 3(b)]. The control signals for the switching chip are brought onto the chip by use of electrical connections [Fig. 3(c)]. The switching chip contains $16 \times 16 \times 16 = 4096$ optical inputs and $16 \times 16 = 256$ optical outputs. As shown in Fig. 3(b), an idle channel can be routed to any of the outputs; this idle channel is connected to the chip electrically. When no inputs of a particular node are selected, the idle channel can provide a time-varying signal that prevents the ac-coupled receivers on the output interface units (see Fig. 2) from chattering.

The control of the individual 16×1 nodes is provided by 17 primary memories and 17 shadow memories, one of each for each of the 16 data inputs and one for the idle channel. Data are loaded into the shadow memories one column at a time (64 columns), and the data are transferred from the shadow to the primary memories during the guard band placed in the data by the input interface unit. There are 23 (26-Mbits/s) electrical inputs to the chip: 20 for the decoded electrical control information, and one each


16, 16×16 switching networks



(a)



(b)



(c)

Fig. 3. Schematic diagram of the optoelectronic VLSI switching chip. (a) Implementation of the network with 16×1 multiplexers or switching nodes showing the physical locations of the modulators and detectors within a switching node. (b) Simplified schematic diagram for a single 16×1 switching node. (c) Circuitry for loading control information into the shadow and primary memories.

for the shift-register clock and shift-register input used to provide enabling inputs for the shadow and primary memories, and one for the idle channel.

The chip was designed and fabricated in a standard $1.0\text{-}\mu\text{m}$ complementary metal-oxide semiconductor (CMOS) process. The 4352 optical detector-modulators were flip-chip bonded to the silicon CMOS circuit, and the substrate removed to allow access to the 850-nm optical ports. Details of the

tion network and the output-packet modules would detect, decode, and resynchronize the inputs and remove the guard band. It would also perform a second translation of the ATM routing information if necessary. An output line unit, located after the output-packet modules, would convert the ATM cells back into the SONET format. A controller interface unit would provide an interface to a PC or workstation that would allow us to program which physical paths through the distribution fabric corresponded to which VPI/VCI values, as well as perform error checking and other diagnostic functions.

For our demonstration, the implementation of one pipe of the distribution network consists of sixteen 16×16 switches, operating at the SONET standard OC-3c rate of 155 Mbits/s. These 16 switches are realized on a single optoelectronic VLSI switching chip with optical inputs and outputs. A fiber-bundle array with 16 input fibers and 16 output fibers accesses four inputs and four outputs from four of the 16 switches on the chip. A digital word generator, controlled by a PC, supplies time-multiplexed data inputs, formatted as the 72-byte cells that the input interface unit would generate from the incoming SONET inputs. The word generator also supplies the out-of-band control signals in place of the path-hunt processor and reconfigures the switch between cells. The experimental system demonstrates that ATM cells could be sent through the distribution network with the appropriate input and output interface units. Alternatively, the switch can be configured as a space switch, with input data originating from any source, including digitized video, as we discuss in Section 7.

3. Switching Chip

All 16 switches for one pipe of the distribution fabric are implemented on one optoelectronic VLSI chip, as shown in block-diagram form in Fig. 3.⁸ The 16×16 networks are implemented with a passive optical fan-out and electronic fan-in by use of 16×1 multiplexers or switching nodes [Figs. 3(a) and 3(b)]. The control signals for the switching chip are brought onto the chip by use of electrical connections [Fig. 3(c)]. The switching chip contains $16 \times 16 \times 16 = 4096$ optical inputs and $16 \times 16 = 256$ optical outputs. As shown in Fig. 3(b), an idle channel can be routed to any of the outputs; this idle channel is connected to the chip electrically. When no inputs of a particular node are selected, the idle channel can provide a time-varying signal that prevents the ac-coupled receivers on the output interface units (see Fig. 2) from chattering.

The control of the individual 16×1 nodes is provided by 17 primary memories and 17 shadow memories, one of each for each of the 16 data inputs and one for the idle channel. Data are loaded into the shadow memories one column at a time (64 columns), and the data are transferred from the shadow to the primary memories during the guard band placed in the data by the input interface unit. There are 23 (26-Mbits/s) electrical inputs to the chip: 20 for the decoded electrical control information, and one each

16, 16×16 switching networks

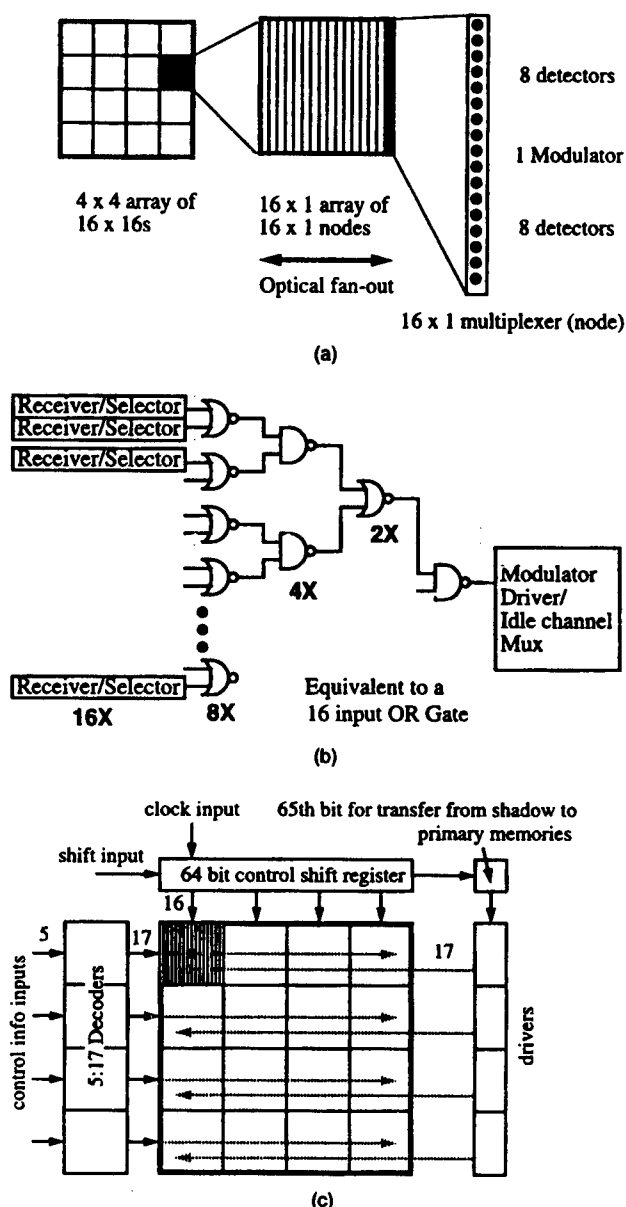



Fig. 3. Schematic diagram of the optoelectronic VLSI switching chip. (a) Implementation of the network with 16×1 multiplexers or switching nodes showing the physical locations of the modulators and detectors within a switching node. (b) Simplified schematic diagram for a single 16×1 switching node. (c) Circuitry for loading control information into the shadow and primary memories.

for the shift-register clock and shift-register input used to provide enabling inputs for the shadow and primary memories, and one for the idle channel.

The chip was designed and fabricated in a standard $1.0\text{-}\mu\text{m}$ complementary metal-oxide semiconductor (CMOS) process. The 4352 optical detector-modulators were flip-chip bonded to the silicon CMOS circuit, and the substrate removed to allow access to the 850-nm optical ports. Details of the



- output fibers (multimode)
- input fibers (single mode)
- read laser fibers (single mode)

Fig. 5. Fiber bundle showing input, output, and read fibers. The fiber bundle accesses part of four 16×16 switches in a column, as illustrated by the four squares in the figure. The input fibers (black) address four of the 16 detectors on each of the nodes within each of the four switches. (A 1×16 phase grating fans the inputs across the switch in the horizontal direction.) Read fibers (gray) provide the read beams that are fanned out to the 16 modulators of a each switch. The output fibers (white) sample four of the 16 outputs (reflected read beams) from the modulators for each switch.

5. Input Lasers

Because all the input lasers and read lasers pass through the same binary phase grating, all lasers must have their wavelengths approximately equal

and stabilized. We chose distributed Bragg reflector lasers at 852 ± 1 nm for these lasers.¹¹ Thermoelectric coolers stabilize the wavelengths for the six lasers that were used to demonstrate a 2×2 section of the chip (including spares) at a trade show at which the ambient temperature was not well controlled. The other 12 lasers do not use thermoelectric coolers. The lasers are pigtailed to 850-nm single-mode fibers, which are connected to the fibers of the bundle by use of standard fiber (ST-type) connectors. The input lasers are electrically driven by a commercial laser-driver chip that accepts emitter-coupled-logic level inputs and provides adjustments for laser bias and peak-current drive levels.

The lasers show excellent wavelength stability. The lasing wavelengths are between 850.8 and 853 nm. The shift in wavelength when modulated at 200 Mbits/s is less than the 0.2-nm resolution of the spectrometer. Wavelength shifts could not be introduced when backreflections were intentionally introduced. Additional measured characteristics of the lasers are summarized in Table 1.

One of the most crucial issues is the temporal waveform under large-signal modulation. Adequate noise margins and pulse-width distortion of the single-ended receivers on the switching chip drove the requirement for a high contrast ratio from the lasers. A contrast ratio of 13 dB (a factor of 20) allows the factor of more than 4 in the ideal case for both the ratio of the detected photocurrent of a high signal to the receiver threshold and the ratio of the receiver threshold to the detected photocurrent of a low signal. These ratios are reduced by variations in optical input power over time, errors in positioning and spot size that cause the coupling into the detector to vary both spatially across the array and over time, and any changes in threshold of the receivers themselves, again both spatially and in time.

Because of the high contrast ratio required, the lasers need to be biased below threshold. Biasing below threshold can lead to an excessive turn-on delay, which also leads to pulse-width distortion in the

Table 1. Measured Characteristics of the Lasers Used in the System Demonstration

Parameter	Symbol	Measured Value
Threshold current	I_{th}	28 mA
Operating current for a given output power P_o	I_{op}	
$P_o = 40$ mW		85 mA
$P_o = 100$ mW		220 mA
Differential slope efficiency	η_d	
$P_o = 5$ mW		0.74 mW/mA
$P_o = 100$ mW		0.4 mW/mA
Wavelength ($P_o = 40$ mW)	λ	852.8 nm
Thermal resistance (junction to ambient)	θ_{ja}	14 °C/W
Divergence		
Parallel to laser junction	$\theta_{ }$	16° FWHM
Perpendicular to laser junction	θ_{\perp}	32° FWHM
Diode series resistance	R_s	5 Ω
Reverse saturation current	I_s	4.5×10^{-16} A
Ideality factor	n	1.7




Fig. 6. Photograph of the system showing the mechanical construction.

system. To keep the turn-on delay below 200 ps, we found that the lasers needed to be biased above 88% of the threshold current. Long-term power stability is an issue when biasing this close to the threshold current. We found qualitatively that the stability is sufficient for operation of the switch at 200 Mbits/s.

6. Physical Design

The mechanical design consists of two parallel plates with pedestals holding the components between plates.⁶ The system is light weight and mechanically symmetric and provides a locking mechanism for components. It is constructed from 6061 aluminum, which has a low density-to-stiffness (ρ/E) ratio, is readily available, has good thermal conductivity, and provides good dimensional stability. The large moment of inertia of the plate-pedestal system provides excellent rigidity but allows convenient access to the optical path for insertion of a view port and optical attenuators during alignment. The position of the elements is held by a locking mechanism consisting of a circular ring that is clamped against the circular components, such as the lenses, fiber bundle, chip mount, and Risley prisms. The chip is mounted in a fixed position and was not rotated during alignment. The fiber bundle and 16×1 grating were rotated to match the rotation of the chip. We achieved chip flatness with respect to the optical axis by mounting the chip with a thin layer of silver paint. Visual inspection showed no change in focus across the chip, which confirms that the chip is mounted reasonably flat.

The mechanical system is mounted in a plastic housing by use of three kinematic contact points. This housing is mounted in a standard electronic equipment frame by use of 1.91-cm rubber spheres as the only vibration isolation. The physical size of the optomechanical system is 14.5 in. \times 3 in. \times 3 in. (36.83 cm \times 7.62 cm \times 7.62 cm). A photograph of the optomechanical system is shown in Fig. 6.

7. Experimental Results

In this section we describe four experiments that demonstrated the operation of the switching demonstrator. First, we operated a section of the switch as a space switch at 200 Mbits/s. Second, we operated that section of the switch as a time-multiplexed switch with ATM-like cells. Then we operated the entire array at 160 Mbits/s by using a larger fan-out grating. Last, we demonstrated digital video switching through a 2×2 section of the switch.

The fiber bundle provides 16 data inputs and four read inputs, interrogating four of the 16 switches within the chip. Because two of the input paths were inaccessible as a result of a permanently latched control bit, 14 lasers were connected to the input fibers. One of the fibers in the bundle is broken; luckily this fiber corresponds to one of the inaccessible channels. These 14 data input lasers were modulated with fixed patterns from the digital word generator. Four lasers were connected to the four read fibers, and these lasers were not modulated. Since there are 16 output fibers and each output fiber can select data from one of four input fibers, there are 64 paths through the network, of which 56 are accessible.

We operated the switch as a slowly reconfigurable space switch, with a data rate of 200 Mbits/s and an electrical-control information rate of 25 Mbits/s during the time that control was being loaded. Laser powers for the individual data input channels were adjusted while the data rate was increased to 300 Mbits/s. At this data rate, all paths were operational immediately following adjustment, although the pulse-width distortion was poor. The paths did not remain operational overnight because of variations in the laser power over time. We believe these variations were caused by the poor electrical connectors that provided a precision voltage reference used in the laser-driver circuit, although temperature-




Fig. 7. Outputs from each of the 16 output fibers, each with four inputs selected at 200 Mbits/s when configured as a space switch. Each of the 16 squares shows the detected outputs from one of the output fibers for the four possible inputs that were supplied. Input lasers were not supplied to fibers 14 and 15 because a permanently latched control bit prevented their selection.

induced threshold changes could also have caused variations in power. At 200 Mbits/s, the switch was not as critically dependent on the laser power, and no measurable change in pulse width was observed for several days.

The data from the output fibers are shown in Fig. 7 for the 64 paths through the network for inputs of a fixed pattern. The bit error rate was measured with a 2^{23} pseudorandom sequence supplied to one of the input lasers, whereas the other lasers had a fixed pattern. Bit error rates below 10^{-13} were observed.

Variations in the optical detected power, both in time and spatially across the array, limit the bit rate of the chip in the system to a rate significantly less than the rate at which one can operate a node within the array. These variations in optical power may cause pulse-width distortion, which, if severe enough, causes bit errors. We have measured⁸ and modeled¹² the pulse-width distortion as a function of both optical power and power-supply voltage for receivers near the four corners of the array. A rough estimate is that a 50% (3-dB) increase or decrease in optical power causes a 0.75-ns increase or a 1.5-ns decrease in the pulse width at 200 Mbits/s.⁸ Also, the variation across the array is nominally less than 400 ps. If we neglect this variation, we can get some indication of the optical detected-power variations across the 64 accessed inputs by comparing the pulse widths shown in Fig. 7 with the measured power dependence of the pulse widths of the individual receivers on the chip [Fig. 7 at 5.6 V (Ref. 8)]. From the measured pulse widths of 4.9 ± 1.1 ns, the calculated detected average powers across the array range from 35–158 μ W. The variations in these powers include not only the difference in laser power, but also differences in the percent of light coupled into the detector win-

dows. Other factors that affect this measurement include pulse-width distortion and skew of the lasers themselves (with respect to one another), variations in threshold across the array, variations in cable delay among cables of a given length, and variations in fiber length. Unfortunately, we could not compare the calculated detected powers with directly measured optical powers because of the problems with the electrical connections to the lasers.

In the second experiment, the digital word generator supplied 576-bit (72-byte) cells at 208 Mbits/s, which is representative of cells that could have been supplied by the input interface unit. This pseudo-cell consisted of an 8-bit preamble (00000101), an 8-bit cell number from 0 to 16 differentially encoded (for example, 01011001 for cell number 2), 69 bytes of data consisting of a repetitive cell representing the input channel number (also differentially encoded), and an 8-bit postamble (0101000). The differentially encoded data was used because it is easier to identify visually compared with nondifferential data with data encoding to reduce long strings of 1's and 0's. The 69 bytes is more than enough for the ATM cell (53 bytes), encoding (<7 bytes), and additional overhead functions. The data repeated every 16 cells.

Figure 8(a) shows one of the detected outputs from one fiber at the end of cell 15 and the beginning of cell 0. The transition between the two cells is quite evident, looking much like an extra bit. This transition did not occur unless the individual multiplexer reconfigured between cells. This rules out current spikes on the power and ground leads during the transfer of data between the primary and shadow memories as the cause of the transition. The likely source of this visible transition is overshoot from the receiver output when it is switched from being disabled to being enabled.

Figure 8(b) shows the outputs from all 16 fibers for the first four cells. Since the fiber bundle could supply only four of the inputs, we chose to cycle the data through these four rather than through all 16. Thus, other than the cell number, the data is repetitive every four cells. Careful inspection showed that the bit patterns have no errors from their expected values (recall that the inputs to fibers 15 and 16 were unavailable).

No special synchronization was performed on the input fibers other than the fixed delay added by the signal generator to account for different-length electrical cables connected between the word generator and the input lasers. The resolution of the delay setting from the word generator was 200 ps. The control signals were delayed relative to the data inputs to set the reconfiguration of the switch in the center of the guard-band interval, as illustrated in Fig. 8. The delay variation across the array was $\sim \pm 1.5$ ns, and all channels reconfigured well within the 40-ns guard band.

The output interface unit would remove the guard band from the data, so that the pulse during the transition would not generate any errors. Since




Fig. 8. Outputs from (a) one channel and (b) 16 channels for the switch operating as a time-multiplexed space switch at 208 Mbits/s. The outputs are shown during reconfiguration from one cell to the next. The data in (b) show the first four time slots of 16.


these units were not built, it was difficult to measure the bit error rate during pseudo-ATM operation. To remove the transition pulse, we summed the detected output from one channel using a resistive splitter-combiner with a corresponding negative pulse from the word generator synchronized to the transition pulse. Another channel of the word generator was programmed with the expected pattern and connected to the bit-error-rate detector, which compares the reference pattern to the detected output (with the removed pulse). Bit error rates below 10^{-13} were measured for this channel in overnight testing.

To operate and diagnose a larger section of the array, we replaced the 1×16 phase grating with a 1×68 grating. A 1×68 grating was used rather than a 1×64 grating because the array was designed in four sections, with an $80\text{-}\mu\text{m}$ space between sections. This grating fanned out each of the 14 inputs across the entire chip in the horizontal direction.


Thus, there were $64 \times 14 = 896$ receivers with incident light, of which 256 were selected, and all 256 optical modulators had incident read beams that provided an output from every node. To access these 256 outputs, we used a video sampling oscilloscope.¹³ A pellicle in the output pupil sampled the light from all 256 modulators and a 40-mm focal length triplet imaged the array onto a high-resolution camera. This lens provided an ideal magnification of the modulator array, nearly fully filling the CCD array. The video sampling oscilloscope uses a pulsed-modulator read beam, synchronized with the modulator drive function (in this case the word generator that drives the data input lasers), but with a slightly different frequency (~ 1 Hz) to strobe or sample the repetitive optical output. The word generator that supplied the data inputs was phase locked to one signal generator set to 160.0000001 MHz, and the four read lasers were connected to pulse generators, which were triggered by a second signal generator set to 20.000000 MHz. The factor of ~ 8 allows one to sample an 8-bit word with a video rate of 1 Hz. The pulse width of the read laser was ~ 1.5 ns, which did cause some rounding of the waveforms, but it was sufficient for determining the extent to which the array was operating.

Figures 9(a) and 9(b) show the array operating under two different control and bias levels. Receivers in alternating columns were designed with different feedback resistors and thus different thresholds or sensitivities. We could not get both sets of receivers to operate concurrently under any bias condition at 160 Mbits/s. It would have been possible to operate both receiver designs concurrently with greater optical powers, at a lower bit rate, or both, because both conditions effectively increase the dynamic range of the receivers. In Fig 9(a), the outputs in the odd-column selected receivers with incident light and valid data and the even-column selected receivers with no incident light. The opposite is true in Fig. 9(b), in which the CMOS supply voltage (V_{dd}) was adjusted (from 5.22 to 4.75 V) to allow operation with the less sensitive receivers. Nearly all channels have recognizable bit patterns.

The measured spread in wavelengths of the input lasers was from 850.8 to 853.0 nm. The design wavelength of the 1×68 grating was 851.7 nm. The differences in wavelength cause the spacing of the spots to change, which, in turn, causes the spots near the edge of the array to be misaligned with respect to the optical windows. The best performance was obtained with wavelengths longer than the design wavelength. This discrepancy could be due to the accuracy of the wavemeter used to measure the wavelength and the degree to which the focal length of the objective lens is known. Four inputs that were fully operational across the array had wavelengths of 852.3–852.8 nm. The difference in wavelength of the lower-wavelength inputs (850.9 nm) from the optimum (852.5 nm) causes a spot shift of $\sim 5\text{ }\mu\text{m}$ near the edge of the array. This is enough to reduce the power coupled into the



(a)



(b)

Fig. 9. Outputs taken by use of the optical sampling spatial oscilloscope from (a) the even and (b) the odd columns from all the nodes in the system at 160 Mbits/s. Because the even and odd columns had different receivers, different voltages were applied for the two figures.

detector by almost 50% near the edge of the array. Also, distortion of the lenses themselves, errors in the positions of the fibers, and aberrations such as field curvature potentially cause even less light to be coupled into these detectors near the edge of the array. Thus, only some of the inputs were opera-

tional across the full field of view. However, all 14 inputs were operational across the 16 nodes that the experiment was originally designed for, when the misalignment for the same wavelength variation was a maximum of $\sim 1.25 \mu\text{m}$.

Before and after these experiments were per-

formed, the system operated for an extended time as a 2×2 video switch. In this mode of operation, video cameras are connected to codecs that encode the video into ATM cells inside a SONET OC-3c frame at 155 Mb/s. The OC-3c optical output from the codecs is fed to a conversion board that converts the 1300-nm multimode signals from the codec to 850-nm single-mode signals required by the switch. After passing through the switch, the detected outputs are converted back to 1300-nm multimode signals and sent to the codecs for conversion back to video, at which the outputs are displayed on television monitors.

This system was displayed at the National Communications Forum Tech Previews for several days and has operated in the laboratory for several months after that. During transportation, the system was subjected to rather harsh vibrations, including being shipped intact in a truck and being dragged on rollers across a parking lot. After transportation back from the show, the system operated with no optical adjustments, qualitatively demonstrating the ruggedness of the optomechanical framework. It is not unusual for an optical system to be subjected to such vibrations and still operate; indeed, robust optomechanical systems are routine for space and avionics applications. However, it is perhaps an important point that systems such as these can be built with robust mechanics, because those in the microelectronics world are skeptical.

8. Conclusion

We have described a new optoelectronic switching system demonstration that implements part of the distribution fabric for a large ATM switch. The system uses a new novel architecture, new device technology, new optical system, and new mechanical system. The system is a major step forward compared with our previous systems, primarily in that the device technology is made by use of high-yield, low-cost, manufacturable VLSI processes, the optical system is drastically simpler, and the optomechanical system is more robust. The technology is advanced enough that terabit optoelectronic switching systems can be contemplated within the next few years.

This work was partially funded by the Defense Advanced Research Projects Agency under Rome Laboratories contract number F 30602-93-C-0166. We thank the authors of Refs. 2 and 8, whose dedicated efforts to build a new device technology enabled this demonstration system to be built. We also

thank D. Vukobratovich for consulting with us on the optomechanical design of the system.

References

1. M. J. Goodwin and A. J. Moseley, "The application of optoelectronic technologies to high-performance electronic processor interconnects," *Opt. Quantum Electron.* **26**, S455-S470 (1994).
2. K. W. Goossen, J. A. Walker, L. A. D'Asaro, S. P. Hui, B. J. Tseng, R. E. Leibenguth, D. Kossives, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller, "GaAs MQW modulators integrated with silicon CMOS," *IEEE Photon. Technol. Lett.* **7**, 360-362 (1995).
3. K. Y. Eng, M. J. Karol, and Y. S. Yeh, "A growable packet switch architecture: design principles and applications," *IEEE Trans. Commun.* **40**, 423-430 (1994).
4. T. J. Cloonan and G. W. Richards, "Terabit per second packet switch having distributed out-of-band control of circuit and packet switching communications," U.S. patent number 5,537,403 (16 July 1996).
5. F. B. McCormick, T. J. Cloonan, A. L. Lentine, J. M. Sasian, R. L. Morrison, R. A. Novotny, M. G. Beckman, S. L. Walker, M. J. Wojcik, S. J. Hinterlong, R. J. Crisci, and H. S. Hinton, "A 5-stage free-space optical switching network with field-effect transistor self-electro-optic-effect smart-pixel arrays," *Appl. Opt.* **33**, 1601-1618 (1994).
6. D. J. Reiley, M. G. Beckman, and J. M. Sasian, "Optomechanical design for free-space optical switching," in *Optoelectronic Packaging*, M. R. Feldman and Y.-C. Lee, eds., *Proc. SPIE* **2691**, 84-90 (1996).
7. T. J. Cloonan, G. W. Richards, and A. L. Lentine, "Optical implementation of a parallel out-of-band controller for large ATM switch applications," in *Optical Interconnects in Broadband Switching Architectures*, T. J. Cloonan, ed., *Proc. SPIE* **2692**, 73-84 (1996).
8. A. L. Lentine, K. W. Goossen, J. A. Walker, L. M. F. Chirovsky, L. A. D'Asaro, S. P. Hui, B. J. Tseng, R. E. Leibenguth, J. E. Cunningham, W. Y. Jan, J. M. Kuo, D. Dahringer, D. Kossives, D. D. Bacon, G. Livescu, R. L. Morrison, R. A. Novotny, and D. B. Buchholz, "Optoelectronic VLSI switching chip with greater than 4000 optical I/O based on flip-chip bonding of GaAs/AlGaAs MQW modulators and detectors to silicon CMOS," *J. Select. Top. Quantum Electron.* **2**, 77-84 (1996).
9. D. J. Reiley and J. M. Sasian, "Optical design of a free-space photonic switching system," *Appl. Opt.* (submitted).
10. J. M. Sasian, R. A. Novotny, M. G. Beckman, and S. J. Hinterlong, "Fabrication of fiber bundle arrays for free-space photonic switching systems," *Opt. Eng.* **33**, 2979-2985 (1994).
11. Spectra Diode Laboratories, Model 5710 (80 Rose Orchard Way, San Jose, Calif., 95134, 1995).
12. R. A. Novotny and A. L. Lentine, "Analysis of pulse-width distortion within an optoelectronic switching system demonstrator," in *Proceedings of the 1996 IEEE/LEOS Summer Topical Meeting on Smart Pixels* (Lasers and Electro-Optics Society, Boston, Mass., 1996), pp. 20-21.
13. R. L. Morrison, S. G. Johnson, A. L. Lentine, and W. Knox, "Design and demonstration of a high-speed multichannel optical oscilloscope," *Appl. Opt.* **36**, 1187-1194 (1996).

GaAs MQW Modulators Integrated with Silicon CMOS

K. W. Goossen, *Member, IEEE*, J. A. Walker, L. A. D'Asaro, *Life Senior Member, IEEE*, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. D. Bacon, D. Dahringer, L. M. F. Chirovsky, A. L. Lentine, *Member, IEEE*, and D. A. B. Miller, *Fellow, IEEE*

Abstract—We demonstrate integration of GaAs-AlGaAs multiple quantum well modulators to silicon CMOS circuitry via flip-chip solder-bonding followed by substrate removal. We obtain 95% device yield for 32×32 arrays of devices with 15 micron solder pads. We show operation of a simple circuit composed of a modulator and a CMOS transistor.

I. INTRODUCTION

FOR many years now a much desired goal of those working on optical interconnects and optical computing has been the integration of high density silicon electronics with high performance GaAs-based optoelectronics. In particular, the possibility of direct optical communication to logic chips has stimulated much work on photonic switching [1]. The most desirable product is one where the silicon circuitry is state-of-the-art, and unaffected by the integration with the optoelectronics. For this reason flip-chip solder bonding to finished silicon chips has been pursued [2]. Furthermore, modulators, which can be fabricated in densities of thousand per chip [3], are the preferred optoelectronic component in many systems such as in [1]. Finally, GaAs-AlGaAs multiple quantum well modulators operating at 850 nm offer the highest performance compared to longer wavelength modulators [4], [5].

In [6], we demonstrated that the GaAs substrate could be removed after flip-chip bonding, allowing operation at 850 nm. This procedure of bonding, followed by substrate removal, has been explored in detail by us, and here we present its application to silicon CMOS, thus fulfilling the above-stated goal. We demonstrate here a 99.9% bond yield with a steadily improving 95% device yield. Furthermore, all aspects of this procedure appear to fit within a manufacturable scheme, with no thin-film handling required as in epitaxial lift-off [7]. We have even demonstrated that completed chips can be sawed without damage, allowing batch fabrication of many chips at once. In [6], the devices operated at high optical intensity (80 kW/cm^2), a huge thermal flux and electrical current density, showing excellent heat-sinking and ohmic contact. The device was thermally cycled from 30°C to




Fig. 1. Three-step hybridization process: 1) Fabrication, aligning, and bonding of modulator chip on silicon chip. 2) Flowing etch-protectant between chips, which is allowed to harden. 3) Removal of GaAs substrate using jet etcher, and deposition of AR coating. The epoxy can be removed after substrate removal, as desired.

100°C over a hundred times, and it showed no degradation, showing the practicality of the technique.

The fabrication procedure is outlined in Fig. 1. Modulators are produced in the GaAs chip whose n and p contacts are coplanar. In [6] this was accomplished by depositing thick gold over the bottom contact. Here we employ implantation [8]. Lead-tin is deposited on these for a solder using photolithography. The silicon chips are obtained from the MOSIS foundry facility. The chip have 1.2 micron linerules. Mating aluminum pads from the modulators are designed on those chips, and a Ti-Pt-Au layer is deposited on them (in our lab) to provide a solder-wettable surface. A precision bonder made by Research Devices in Piscataway, NJ was employed to bond the chips together. Two-micron accuracy is routine.

A key feature of the technique for flip-chip bonding then substrate removal is the etching of outer mesas around the devices into the substrate. Then, when the substrate is removed by applying a chemical stream to it (that stops on the AlGaAs stop-etch layer), isolated devices will be left. This is desirable since, if the stop etch layer was left extending over the whole chip, slight warpages would cause it to break, possibly damaging the modulators. This procedure requires placing something between the mesas so that the substrate etchant does not attack the front faces of the chips. The substrate etchant, $100:1 \text{ H}_2\text{O}_2:\text{NH}_4\text{OH}$, does not attack Si or Al appreciably.

Manuscript received November 1, 1994; revised December 9, 1994.

K. W. Goossen, J. A. Walker, and D. A. B. Miller are with AT&T Bell Laboratories, Holmdel, NJ 07733 USA.

L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. Dahringer, and L. M. F. Chirovsky are with AT&T Bell Laboratories, Murray Hill, NJ 07974 USA.

A. L. Lentine is with AT&T Bell Laboratories, Naperville, IL 60566 USA. IEEE Log Number 9409250.




Fig. 2. Photo of integrated GaAs modulators on silicon CMOS. Results on the transistor modulator circuit (near bottom of photo) are presented here. Results on the complex circuits will be reported later.

However, it would attack the GaAs regions of the modulators. To protect the front faces of the chips, a silica-filled epoxy was flowed between the chips and allowed to harden, as shown in the middle pictorial of Fig. 1. This was done by depositing a bead of the epoxy on the side of the GaAs substrate using a optical fiber manipulated by a precision stage. The epoxy then wicked neatly between the chips. The chip is heated to 100 °C to reduce the viscosity of the epoxy so that it flows between the chips more easily. It is possible to meter the amount of epoxy so that it just fills the volume between chips. The epoxy is then cured by baking the chip at 100 °C for one hour. Epoxies have been used previously in this manner in flip-chip bonded assemblies to provide hermetic sealing and increase robustness [9]. For those applications the epoxy is termed an encapsulant, or underfill. Here we call it an interchip flowable hardener, to express the added function of providing a surface between the chips that is impenetrable by the substrate etch. The epoxy can be removed after substrate removal by applying a dry plasma etch using 5:1 O₂:CF₄ flow rates.

In these devices, a Ti-Au pad, place next to the *n* ohmic contact, is used as an integral reflector. We have previously demonstrated that modulators such as these using pure Au pads have performance equal to the best monolithic GaAs modulators [10]. Here the Ti was added to provide better sticking of the Au and so improve yield. Unfortunately our Ti-Au only has about 40% reflectivity, so the modulators here have marginal performance. We are developing schemes to use pure Au reflectors with good adhesion.

We have fabricated CMOS chips with switching node electronics (Fig. 2). Results on the switching nodes will be discussed in a later paper. Here we discuss device performance, consisting of three tests: *n*-ohmic bond test arrays, LED device



Fig. 3. GaAs modulator test array on silicon.




Fig. 4. Best section of LED (forward-biased modulator) array with 98/99 working devices.

test arrays (forward-biased modulators), and simple circuits (near bottom, Fig. 2; inset, Fig. 6). Our *n*-ohmic bond testers consisted of daisy-chains of devices with only *n*-contacts. For these we obtained 99.94% bond yield for 15×15 micron solder pads (Table I). However, our LED test arrays had only 95% device yield (Fig. 4). We have attributed this to an observable intermetallic reaction that occurs between the solder and the *p*-type metal during solder reflow (melting), which is shown in Fig. 6. This reaction is visible in about half the devices




Fig. 5. Two bonded devices viewed through substrate with infrared microscope. The "tab"-shaped metal is the p -ohmic contact. The device on the left shows no degradation. The device on the right shows a reaction with the solder. There is a strong correlation with this observation and a failed device. We are examining methods of bonding without reflow to avoid this effect.




Fig. 6. Reflectivity of modulator in inset circuit versus gate-source voltage, showing electrical integration. Modulation is degraded compared to earlier devices using Au as a reflector (here Ti-Au is used).

with an infrared microscope. We have measured that if an LED is dark, there is a 96% probability that it also exhibits the intermetallic reaction. The reaction could be avoided by not performing reflow. However, it is during reflow, which is performed in a solder flux, when the solder oxide is removed. We have attempted bonding without reflow, by subjecting the chips to a plasma before bonding to remove the oxide. We have obtained sections of arrays as large as 12×42 with uniform illumination of all devices, but the results are still incomplete.

Finally, we show here a simple CMOS-modulator circuit (inset, Fig. 6). This circuit is shown on the bottom of the photo in Fig. 2. By charging the gate of the transistor, the

TABLE I
YIELD OF ARRAYS OF GaAs DEVICES WITH TWO
 n -OHMIC CONTACTS SOLDER-BONDED TO SILICON

pad size (microns)	array-size	bond yield
10x10	32x80	99.88 %
15x15	28x60	99.94 %
20x20	24x50	100 %

transistor turns on and the modulator is biased. In Fig. 6 we show the turn-on characteristic. The design gate threshold of this transistor is about one volt. The turn-on of the modulator at 200 mV is consistent with this since the modulator had only nanowatts of optical power on it, so required only subthreshold operation of the transistor.

II. CONCLUSION

We have demonstrated a practical method of integrating GaAs modulators onto silicon circuits via flip-chip bonding, followed by substrate removal. We obtain 95% device yield, and indicate that this can improve to 99.9%. We have demonstrated a simple transistor-modulator circuit to prove viability. More complex circuits will be reported at a later date.

REFERENCES

- [1] A. L. Lentine, R. A. Novotny, T. J. Cloonan, L. M. F. Chirovsky, L. A. D'Asaro, G. Livescu, S. Hui, M. W. Focht, J. M. Freund, G. D. Guth, R. E. Leibenguth, K. G. Glogovsky, and T. K. Woodward, "4 x 4 arrays of FET-SEED embedded control 2 x 1 optoelectronic switching nodes with electrical fan-out," *IEEE Photon. Technol. Lett.*, vol. 6, p. 1126, 1994.
- [2] J. Wieland, H. Melchior, M. Q. Kearley, C. Morris, A. J. Moseley, M. G. Goodwin, and R. C. Goodfellow, "Optical receiver array in silicon bipolar technology with self-aligned, low parasitic III/V detectors for DC-1 Gbit/s parallel links," *Electron. Lett.*, vol. 27, p. 2211, 1991.
- [3] A. L. Lentine, F. B. McCormick, R. A. Novotny, L. M. F. Chirovsky, L. A. D'Asaro, R. F. Kopf, J. M. Kuo, and G. D. Boyd, "A two kbit array of symmetric self-electro-optic effect devices," *IEEE Photon. Technol. Lett.*, vol. 2, p. 51, 1990.
- [4] K. W. Goossen, M. B. Santos, J. E. Cunningham, and W. Y. Jan, "Independence of absorption coefficient-linewidth product to material system for multiple quantum wells with excitons from 850 nm to 1064 nm," *IEEE Photon. Technol. Lett.*, vol. 5, p. 1392, 1993.
- [5] R. N. Pathak, K. W. Goossen, J. E. Cunningham, and W. Y. Jan, "In-GaAs-InP p -i(MQW)- n surface normal electroabsorption modulators exhibiting better than 8:1 contrast ratio for 1.55 μ m applications grown by gas-source MBE," *IEEE Photon. Technol. Lett.*, vol. 6, no. 12, pp. 1439-1441, 1994.
- [6] K. W. Goossen, J. E. Cunningham, and W. Y. Jan, "GaAs 850 nm modulators solder-bonded to silicon," *IEEE Photon. Technol. Lett.*, vol. 5, p. 776, 1993.
- [7] C. Camperi-Ginestet, M. Hargis, N. Jokerst, and M. Allen, "Alignable epitaxial liftoff of GaAs materials with selective deposition using polyimide diaphragms," *IEEE Photon. Technol. Lett.*, vol. 3, p. 1123, 1991.
- [8] L. A. D'Asaro, L. M. F. Chirovsky, E. J. Laskowski, S. S. Pei, T. K. Woodward, A. L. Lentine, R. E. Leibenguth, M. W. Focht, J. M. Freund, G. G. Guth, and L. E. Smith, "Batch fabrication and operation of GaAs-AlGaAs field-effect transistor-self-electro-optic effect device (FET-SEED) smart pixel arrays," *IEEE J. Quantum Electron.*, vol. 29, p. 670, 1993.
- [9] G. O'Malley, J. Giesler, and S. Machuga, "The importance of material selection for flip chip on board assemblies," *IEEE Trans. Components, Hybrids, Manufacturing Technol.*, part B, vol. 17, p. 248, 1994.
- [10] K. W. Goossen, "Integration of GaAs modulators to silicon IC's by monolithic and hybrid techniques," in *Proc. 1994 IEEE Sarnoff Symposium*, section IIc.

To put this condition into perspective, integration of the right hand side of eqn. 4 along the waveguide leads to the curves shown in Fig. 3. These curves give the minimum permissible distance δz along the waveguide as a function of the core increase δn for different values of the waveguide parameters. For all curves we have $\lambda = 1.3$, $\rho = 3.68\mu\text{m}$ and $n_{cl} = 1.447$. The nominal core index (before UV-exposure) is chosen such that the waveguide V -values are 1.5, 1.75 and 2.0 as indicated. In practice, the distance required for a given index increase should be much larger than that predicted by each curve. For example, if an increase in index of $\delta n = 0.002$ is planned, then the length over which this increase occurs should be significantly larger than $100\mu\text{m}$ for a waveguide with $V = 2.0$.

Conclusions: We have proposed a new technique for reducing the minimum bend radius in singlemode planar waveguide circuits, which relies only on the photosensitivity of the core material.

Acknowledgment: R. Jarvis is the recipient of an Australian Postgraduate Award (Industry) in collaboration with Siemens Ltd, Melbourne.

© IEE 1997

Electronics Letters Online No: 19970609

24 March 1997

R.A. Jarvis, J.D. Love and F. Ladouceur (Australian Photonics Cooperative Research Centre, Optical Sciences Centre, The Australian National University, Canberra, ACT 0200, Australia)

References

- SHIMA, T., SHIRAIISHI, K., and KAWAKAMI, S.: 'Waveguide-bend configuration with low-loss characteristics', *Opt. Lett.*, 1986, 11, pp. 736-738
- PENNINGS, E.C.M.: 'Bends in optical ridge waveguides (modeling and experiments)', PhD Thesis, Technische Universiteit Delft, 1990
- LADOUCEUR, F., and LOVE, J.D.: 'Silica based buried channel waveguides and devices' (Chapman & Hall, London, 1995)
- MAXWELL, G.D., KASHYAP, R., AINSLIE, B.J., WILLIAMS, D.L., and ARMITAGE, J.R.: 'UV written $1.5\mu\text{m}$ reflection filters in singlemode planar silica guides', *Electron. Lett.*, 1992, 28, pp. 2106-2107
- BAZYLENKO, M., GROSS, M., CHU, P.L., and MOSS, D.: 'Photosensitivity of Ge-doped silica deposited by hollow cathode PECVD', *Electron. Lett.*, 1996, 32, pp. 1198-1199
- LOVE, J.D.: 'Application of low-loss criterion to optical waveguides and devices', *IEE Proc. Optoelectron. J.*, 1989, 136, pp. 225-228

Optoelectronic VLSI switching chip with greater than 1Tbit/s potential optical I/O bandwidth

A.L. Lentine, K.W. Goossen, J.A. Walker, J.E. Cunningham, W.Y. Jan, T.K. Woodward, A.V. Krishnamoorthy, B.J. Tseng, S.P. Hui, R.E. Leibenguth, L.M.F. Chirovsky, R.A. Novotny, D.B. Buchholz and R.L. Morrison

Indexing terms: Optical interconnections, Optoelectronic devices, VLSI

The authors describe an optoelectronic switching chip with 1024 differential optical inputs and 1024 differential optical outputs with individual channels tested above 600 Mbit/s. The technology for the chip consists of 850nm GaAs/AlGaAs multi-quantum well (MQW) detectors and modulators, flip-chip bonded onto silicon CMOS with substrate removal to allow access to the optical devices.

Recently, a VLSI optoelectronic chip [1] and system [2] have been demonstrated, implementing part of a simplified distribution fabric for a growable packet ATM switch [3]. That chip, designed in $1\mu\text{m}$ CMOS, contained 256 16×1 nodes (or 16 16×16 switches with optical fan-out) operating at a maximum speed of 450 Mbit/s. The system operated at 208 Mbit/s as a time multiplexed switch, capa-

ble of routing ATM cells at OC-3c rates (155 Mbit/s) with an appropriate out-of-band controller. The chip described here, designed in $0.8\mu\text{m}$ CMOS, contains 64 16×16 switches (Fig. 1). The 16×16 switches are implemented by fanning out the electrical outputs from 16 differential receivers to 16 16×1 multiplexers, each with a differential optical output. Control of the chip is electronic. The combination of increased density ($\lambda = 0.4\mu\text{m}$ against $\lambda = 0.5\mu\text{m}$ for λ -based design rules), the use of a third level metal with circuitry underneath the flip-chip bonding pads [4], and electrical fan-out allows four times the functional circuitry in the same area. The new circuit contains over 450K FETs compared to 160K for the old circuit. The chip uses the same MQW diode arrangement as our previous chip and was mounted in the same package. The MQW diodes are bonded with $15\mu\text{m}$ square bumps, have $\sim 11\mu\text{m}$ active areas within the optical windows, are on $80\mu\text{m}$ centres and are arranged in an array of 64×68 diodes.




Fig. 1 Optoelectronic switching chip designed in $0.8\mu\text{m}$ CMOS

- a Block diagram of chip: each square represents a 16×16 switch
b Each switch implemented using 16 16×1 multiplexers with electrical fan-out of receiver outputs (arrows)
c Optical differential receivers (light shading) and transmitters (dark shading) that are overlaid on top of switches; receivers are mapped by number to horizontal fan-out lines (arrows) and transmitters are mapped from multiplexers (column numbers)

The receiver (Fig. 2) uses a modified design of that described in [5]. The sensitivity is not quite as good as that in [5], partially because an imbalance was intentionally introduced at the stage following the receiver and before a large electrical fan-out driver to reduce the static dissipation. Also, the receiver outputs from the neighbouring 16×16 switches were inadvertently pair-wise connected and testing was carried out by driving both receivers in tandem using a 1×2 binary phase grating. Unequal optical powers on the inadvertently connected receivers may have caused signal distortion in the common line. Our previous chip with single ended receivers showed good uniformity ($< \pm 400\text{ps}$ gate delay variations) across the chip [1]. We expect the variations in these circuits to be either comparable or less.




Fig. 2 Sensitivity against bitrate for switching nodes

The large number of active receivers leads to a static power dissipation of almost 5W. The exciton shift due to thermal heating of the circuit was found to be $\sim 10\text{nm}$ at both at the centre and the edges of the array. Using a shift of $0.28\text{nm}/^\circ\text{C}$, the thermal resistance of the package was found to be $7^\circ\text{C}/\text{W}$. The uniformity of the position of the exciton peak in reflectivity while dissipating 5W indicates the temperature uniformity across the array is within

a few degrees. The chip mount was specifically designed for temperature uniformity [6].

There are 16 384 (64×16×16) paths through the switching chip. 16 paths were measured at a data rate of 625Mbit/s (Fig. 3), sampling the upper left 16×16 switch, exercising each receiver and each modulator driver from that switch. Owing to the architecture of the chip, we would not expect large delay variations between inputs, this was indeed the case on our earlier chip that was more thoroughly characterised [1]. Individual channels had bit error rates below 10⁻¹¹ at 600Mbit/s and below 10⁻⁹ at 625Mbit/s. While all paths were not measured, all but seven MQW diodes luminesced under forward bias. Thus, if there are no additional problems in the silicon circuit, nearly all paths should be operational.




Fig. 3 16 outputs from upper left 16×16 switch

Output *i* selects input *i* at 625Mbit/s with bit pattern of 01000111. Zero line is accurate, i.e. the contrast ratio is approximately 2:1

We have described initial characteristics of an optoelectronic switching chip with a potentially > 1Tbit/s I/O bandwidth. The chip has one simple design error that may be limiting its performance and is easily corrected. While crosstalk, thermal spatial variations under dynamic operation, and delay variations need to be rigorously characterised before this I/O bandwidth can be used in a system, this chip further illustrates the potential of hybrid optoelectronic VLSI smart pixel technologies.

© IEE 1997

13 February 1997

Electronics Letters Online No: 19970567

A.L. Lentine, R.A. Novotny, D.B. Buchholz and R.L. Morrison (Bell Laboratories, Lucent Technologies, Naperville, IL 60566, USA)

K.W. Goossen, J.A. Walker, J.E. Cunningham, W.Y. Jan, T.K. Woodward and A.V. Krishnamoorthy (Bell Laboratories, Lucent Technologies, Holmdel, NJ 07733, USA)

B.J. Tseng, S.P. Hui and L.M.F. Chirovsky (Bell Laboratories, Lucent Technologies, Murray Hill, NJ 07974, USA)

R.E. Leibenguth (Bell Laboratories, Lucent Technologies, Breinigsville, PA 18031, USA)

References

- LENTINE, A.L., GOOSSEN, K.W., WALKER, J.A., CHIROVSKY, L.M.F., D'ASARO, L.A., HUI, S.P., TSENG, B.J., LEIBENGUTH, R.E., CUNNINGHAM, J.E., JAN, W.Y., KUO, J.M., DAHRINGER, D., KOSSIVES, D., BACON, D.D., LIVESCU, G., MORRISON, R.L., NOVOTNY, R.A., and BUCHHOLZ, D.B.: 'Optoelectronic VLSI switching chip with greater than 4000 optical I/O based on flip chip bonding of GaAs/AlGaAs MQW modulators and detectors to silicon CMOS', *J. Sel. Top. Quantum Electron.*, 1996, 2, pp. 77-84
- LENTINE, A.L., REILEY, D.J., NOVOTNY, R.A., MORRISON, R.L., SASIAN, J.M., BECKMAN, M.G., BUCHHOLZ, D.B., HINTERLONG, S.J., CLOONAN, T.J., RICHARDS, G.W., and MCCORMICK, F.B.: 'ATM distribution network using an optoelectronic VLSI switching chip', *Appl. Opt.*, (to be published)

- CLOONAN, T.J., and RICHARDS, G.W.: US Patent number 5,537,403
- KRISHNAMOORTHY, A.V., LENTINE, A.L., GOOSSEN, K.W., WALKER, J.A., WOODWARD, T.K., FORD, J.E., APLIN, G.F., D'ASARO, L.A., HUI, S.P., TSENG, B., LEIBENGUTH, R.E., KOSSIVES, D., DAHRINGER, D., CHIROVSKY, L.M.F., and MILLER, D.A.B.: '3-D integration of MQW SEED modulators over active sub-micron CMOS circuits: 375Mbit/s transimpedance receiver-transmitter circuit', *IEEE Photonics Technol. Lett.*, 1995, 7, pp. 1288-1290
- WOODWARD, T.K., KRISHNAMOORTHY, A.V., LENTINE, A.L., GOOSSEN, K.W., WALKER, J.A., CUNNINGHAM, J.E., JAN, W., D'ASARO, L.A., CHIROVSKY, L.M.F., HUI, S.P., TSENG, B., KOSSIVES, D., DAHRINGER, D., and LEIBENGUTH, R.E.: '1Gbit/s two-beam transimpedance smart pixels optical receivers made from hybrid GaAs MQW modulators bonded to 0.8μm silicon CMOS', *IEEE Photonics Technol. Lett.*, 1996, 7, pp. 422-424
- BUCHHOLZ, D.B., LENTINE, A.L., and NOVOTNY, R.A.: 'Thermal considerations in the design of optoelectronic device mounts', *Topical Meeting on Photonics in Switching*, 1995, pp. 118-120

Generalised mathematical model for raindrop size distribution (RSD) for application in radiowave propagation and meteorological studies

K.I. Timothy and S.K. Sarkar

Indexing terms: Rain, Radiowave propagation

A generalised mathematical model for the raindrop size distribution (RSD) is presented. The model RSD is compared with the observed RSD over some tropical stations. The rain attenuation, computed from the model is also discussed.

Introduction: Raindrop size distribution (RSD) over different climatic conditions and in various types of rainfall is of vital importance in radiowave propagation studies. Doppler radar data interpretation and the computation of certain parameters in cloud physics. Estimation of the effects of rain on radiowave propagation involves three steps, namely (i) computation of scattering amplitudes, (ii) evaluation of farfield intensities and (iii) assessment of overall signal degradation. Unfortunately, meteorological uncertainties limit theoretical studies. The limitation in computation of scattering amplitudes arises from the uncertainty in the shape of the raindrop. Similarly, the RSD and spatial distribution of rainfall over the path length are the two uncertainties involved in the second and third steps, respectively. However, considerable improvements have already been in determining the shape of the raindrops, as well as the spatial distribution of rainfall over the path length [1, 2]. But the information available on RSD is still far from completion. A marked difference in RSD between tropical and temperate climates has been reported many times. The Marshall and Palmer RSD model is being widely used in radiowave propagation studies, while meteorologists prefer log-normal and gamma RSD models.

This Letter aims at proposing a generalised mathematical RSD model which can represent the RSD for all types of rainfall over tropical climates.

Mathematical model: The well-known integral equation that relates RSD or number density $N(D)$, terminal velocity $V(D)$, and the drop diameter (D) with the rainfall rate (R) is given by

$$R = 6\pi \times 10^{-4} \int D^3 \times N(D) \times V(D) dD \quad (1)$$

where R is mm/h, $N(D)$ is in m⁻³mm⁻¹, $V(D)$ is in m/s and D is in mm. Since this integral equation is well accepted and is being used for various calculations in related fields, any RSD model (if it has to be accepted) should not violate its consistency.

The observation on $V(D)$ which are generally considered most accurate, are those by Gunn and Kinger [3]. The best fit curve for their data can be represented by the following equation. Here, the maximum diameter of a raindrop is considered to be equal to 6mm

High-Speed Optoelectronic VLSI Switching Chip with >4000 Optical I/O Based on Flip-Chip Bonding of MQW Modulators and Detectors to Silicon CMOS

Anthony L. Lentine, *Member, IEEE*, Keith W. Goossen, *Member, IEEE*, J. A. Walker, Leo M. F. Chirovsky, *Member, IEEE*, L. Arthur D'Asaro, *Senior Life Member, IEEE*, S. P. Hui, B. J. Tseng, R. E. Leibenguth, J. E. Cunningham, W. Y. Jan, Jen-Ming Kuo, *Member, IEEE*, D. W. Dahringer, D. P. Kossives, D. D. Bacon, Gabriela Livescu, *Member, IEEE*, R. L. Morrison, Robert A. Novotny, *Member, IEEE*, and D. B. Buchholz

Abstract—We present the first high-speed optoelectronic very large scale integrated circuit (VLSI) switching chip using III-V optical modulators and detectors flip-chip bonded to silicon CMOS. The circuit, which consists of an array of 16×1 switching nodes, has 4096 optical detectors and 256 optical modulators and over 140K transistors. All but two of the 4352 multiple-quantum-well diodes generate photocurrent in response to light. Switching nodes have been tested at data rates above 400 Mb/s per channel, the delay variation across the chip is less than ± 400 ps, and crosstalk from neighboring nodes is more than 45 dB below the desired signal. This circuit demonstrates the ability of this hybrid device technology to provide large numbers of high-speed optical I/O with complex electrical circuitry.

I. INTRODUCTION

THE use of optical interconnections between electronic devices is expected to alleviate the communication bottleneck that exists in today's large electronic switching and computing systems. A particularly attractive approach is to integrate the optical detectors and modulators or sources directly onto electronic circuitry. Up to now, circuits such as these have tended to be limited in speed, such as the case with ferroelectric liquid-crystal devices [1], [2], or complexity, such as the monolithic field-effect transistor self-electrooptic effect devices (FET-SEED's) [3]. The hybrid integration of multiple-quantum-well (MQW) modulators with silicon CMOS offers the potential for both high speed and high complexity. Circuits made using hybrid bump-bonding techniques [4]–[10] have achieved numbers as high as 16-K optical I/O in a spatial light modulator [7] and data rates above 1 Gb/s in receiver transmitter pairs [8].

Manuscript received March 15, 1996; revised July 17, 1996. This work was sponsored in part by ARPA under Air Force Rome Laboratories Contract F 30602-93-C-0166.

A. L. Lentine, R. L. Morrison, R. A. Novotny, and D. B. Buchholz are with Bell Laboratories, Lucent Technologies, Naperville, IL 60566 USA.

K. W. Goossen, J. A. Walker, J. E. Cunningham, and W. Y. Jan are with Bell Laboratories, Lucent Technologies, Holmdel, NJ 07733 USA.

L. M. F. Chirovsky, L. A. D'Asaro, S. P. Hui, B. J. Tseng, J.-M. Kuo, D. W. Dahringer, D. P. Kossives, D. D. Bacon, and G. Livescu are with Bell Laboratories, Lucent Technologies, Murray Hill, NJ 07974 USA.

R. E. Leibenguth is with Bell Laboratories, Lucent Technologies, Breinigsville, PA 18031 USA.

Publisher Item Identifier S 1077-260X(96)07973-7.




Fig. 1. Block diagram of an optoelectronic chip. Control information, clock, and shift inputs are electrical.

We present the first high-speed optoelectronic VLSI switching chip using this technology. The circuit implements sixteen 16×16 fully interconnected switches, where each of these switches contains sixteen 16×1 multiplexers or switching nodes with external optical fan-out and electrical fan-in as shown in the block diagram in Fig. 1. The chip implements part of an optoelectronic distribution network for an ATM switching demonstration [11]. The circuit has 4096 optical detectors and 256 optical modulators. The control information is brought into the chip via electrical connections. It is decoded and routed to the individual nodes in a column-by-column basis using the parallel outputs from a shift register.

Fig. 2 shows a schematic of a 16×1 switching node consisting of 16 receiver/selectors, an OR tree, control memories, and an output section. Throughout the rest of this paper, we will use the term "switching node" rather than the term "multiplexer" because a node contains elements, such as receivers, control memories, and modulator drivers, not normally associated with a multiplexer. The receiver/selectors serve two functions. First, they convert the 16 optical inputs into electrical signals. Second, based on information stored locally in control memories, they "select" which of the 16 inputs is to be routed to the output. That is, only one of the receiver/selectors will be enabled at a time. The outputs from the receiver/selectors are then routed to a 16-input OR gate tree, implemented with four stages of two-input NAND/NOR logic and then routed to the modulator driver section. Using a fan-in of two in these gates

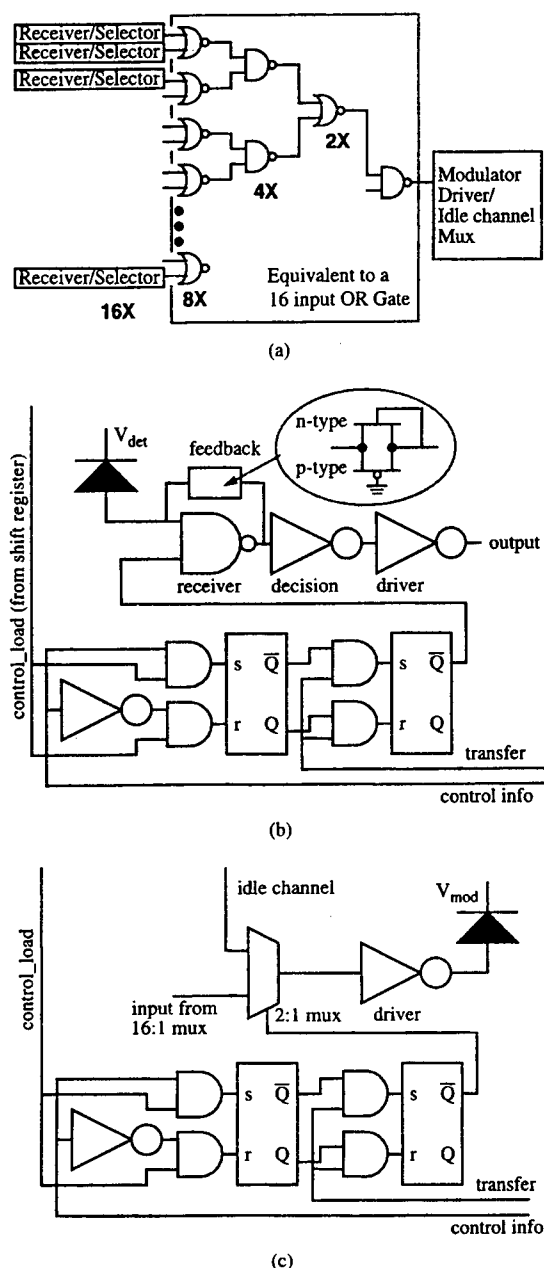


Fig. 2. Schematic diagrams of a 16×1 node: (a) overall schematic, (b) receiver/selector schematic, and (c) modulator driver/idle channel mux schematic.

minimizes capacitive loading because the gates are spread out in space across the node. The longest electrical trace is approximately $320 \mu\text{m}$, with an estimated capacitance of 25 fF. In the output section, the data passes through a 2×1 multiplexer and then to a final modulator-driver inverter. If none of the receivers are selected, the 2×1 multiplexer inserts an idle signal, which is routed onto the chip electrically.

Each individual optical input of the node has a shadow and a primary memory associated with it that determine if that particular optical input is the one of sixteen that is to be routed to the node output as shown in Fig. 2(b). There is also a shadow and primary memory associated with the

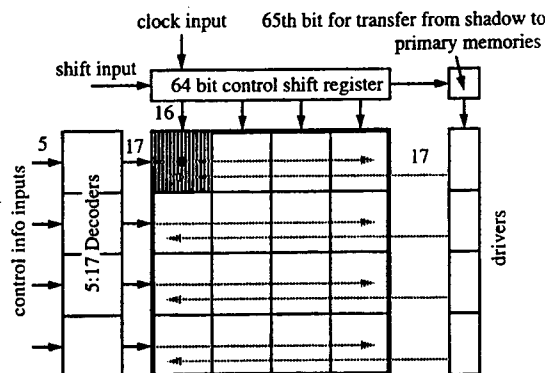


Fig. 3. Block diagram of the electronic control circuitry of the chip. All control inputs are electronic. The control information is sequentially loaded into the shadow memories, one column at a time, using the outputs from each bit of the shift register to enable writing of the memories. A signal to transfer the data from the shadow to primary memories is derived from the 65th bit of the shift register.

output section that determines whether the idle channel is selected as shown in Fig. 2(c). The shadow memories are loaded while data is passing through the nodes and then the control information is transferred from the shadow to the primary memories during a guard band in the data.

The control information is read into the shadow memories in a column-by-column basis as illustrated in Fig. 3. Four 5:17 decoders on the left side of the array provide the control information which is routed horizontally to the nodes in each of the four rows. The five input bits per row provide control information for the 16-receiver/selectors plus one bit for the idle channel. A shift register provides an input to each column of nodes that enables writing of the shadow memories with the control information bits, one column of nodes at a time. The 65th bit of the shift register provides a signal that transfers data from the shadow to the primary memories. Thus, the signal input to the shift register is a logic one followed by a string of logic zeros. The input data and control load signals are synchronized in time by a common control load clock and master-slave flip-flops in the control information inputs. Clock frequencies above 100 Mb/s have been used to load the control information into the switching nodes. This corresponds to a reconfiguration time of ~ 655 ns to load control information into the shadow memories in each of the 64 columns of nodes in the array and transfer this control information from the shadow to primary memories. The time required for the receivers to become active after this transfer step was measured to be less than 5 ns.

Each electrical control bit is connected to one of two inputs to an electrical differential amplifier. A reference voltage is connected to the other side of the amplifier. The amplifiers unique design [12] enabled the threshold to be set anywhere from 0.5 to 4.5 V (with a 5-V supply), whereas standard differential amplifiers often must have a reference near the center of the voltage range. For compatibility with standard electronic circuitry, the reference voltage can be nominally set to 3.7 V, the threshold of positive emitter coupled logic (P-ECL). However, almost all of the testing was done with the reference voltage at 0.5 V, with the input voltage swing from 0 to 1 V.

Each of the 4096 optical receivers is a dc coupled transimpedance design with a novel nonlinear feedback element to improve the dynamic range [9]. This feedback element consists of a parallel combination of a p-type FET with its source grounded that acts as a resistor for small photocurrents and an n-type FET with its source connected to its drain that acts as a voltage clamp that limits the voltage swing for larger photocurrents. Because only one of 16 inputs is routed to the node output, the receiver resembles a two-input NAND gate. One input to the NAND gate is the detected photocurrent and the other input is the signal from the control memory that determines, based upon the control information, whether that input is the selected one. Performing the selection process in this way has the important advantage of reducing the static dissipation in the unselected receivers. Because each active receiver dissipates ~ 2.5 mW, this reduces the static dissipation of the chip from ~ 10 W if all receivers were continuously biased to less than 1 W when only 256 are selected. The dissipation was determined from SPICE simulations.

The fabrication is similar to previous circuits that we've made using this technology [6], but this circuit was a challenge in that the number of optical modulators/detectors was much larger than the earlier arrays (4392 versus 256 [10]). The silicon CMOS circuit was made using standard foundry silicon with two layers of interconnect metallization and was designed with minimum linewidths and gate lengths of $1\ \mu\text{m}$. Barrier metals and solder are deposited on large sections of the silicon VLSI wafers at once, and, while it should be possible to deposit these metals on entire wafers, this was not done because of the risk of destroying an entire wafer of circuits with an error on this step. The optical modulators and detectors have 95 periods of GaAs-Al_{0.3}Ga_{0.7}As quantum wells with $\lambda_0 = 844$ nm designed for operation at 8 V at 850 nm. There is no epitaxially grown bragg reflector, because the top layer of metallization on the GaAs devices acts, after flip-chip bonding, as the optical mirror for these reflection mode devices. An additional stop etch layer protects the bottom side of the diodes from the substrate etchant. The modulator/detector arrays are made using similar process steps to our earlier monolithic arrays [3] except for a deep etch between individual p-i-n diode mesas that mechanically isolates the diodes after substrate removal. The silicon and GaAs chips are bonded using a precision bonder, and epoxy is introduced into the $\sim 3\text{-}\mu\text{m}$ gap between the two chips. The epoxy protects the surface of the GaAs-AlGaAs modulator array from the substrate etchant and provides additional physical support of the modulator array. After flip-chip bonding, the GaAs substrate is removed with a selective etch to provide optical access to and mechanical isolation of the individual diodes. Mechanically isolating the diodes greatly reduces the strain between the isolated GaAs diodes and the silicon chip compared to devices that have a transparent substrate that remains intact. Lastly, the device is packaged and antireflection-coated. Unlike earlier arrays [10], this array has no missing modulators near the edges of the array. The primary reason for this is the use of an InGaP stop etch layer [13], which has a greater resistance to the substrate etchant, than the AlGaAs in the previous design. A diagram of a cutaway view of the device is shown in Fig. 4(a) and a





Fig. 4. (a) Pictorial cutaway view of flip-chip bonded MQW modulators and detectors on silicon VLSI and (b) photograph of a section of the array. Rectangles are the optical modulators and detectors which are $\sim 23\text{-}\mu\text{m} \times 53\text{-}\mu\text{m}$ and are located on $80\text{-}\mu\text{m}$ centers.

photograph of a section of the array after substrate removal is shown in Fig. 4(b).

The metallic pads for bump bonding are $15\text{-}\mu\text{m} \times 15\text{-}\mu\text{m}$ with a $15\text{-}\mu\text{m}$ space between the n-type and p-type connections of a diode. The individual diodes are on $80\text{-}\mu\text{m}$ centers. The active area of the optical window is $11\text{-}\mu\text{m} \times 11\text{-}\mu\text{m}$, which is reduced from the pad and diode size by an isolation implant. No circuitry was placed underneath the bump-bond pads in this design, even though we have now made circuits with FET's underneath the pads [9]. In the vertical direction, the inputs to a particular node are arranged in a column of 17 diodes consisting of eight detectors, the output modulator, and eight detectors as shown in Fig. 1. Since there are four nodes in the vertical direction, there are 68 MQW diodes down a column. In the horizontal direction, there are 64 nodes, but they are arranged in four groups of 16 with a gap of $80\ \mu\text{m}$ between groups. The gap provides space for V_{dd} and ground power supply connections, so that the array can be powered in sections and voltage variations on the power supply leads are minimized. The optical field of view is $67 \times 80\text{-}\mu\text{m}$ or 5.36 mm in the horizontal direction and $68 \times 80\text{-}\mu\text{m}$ or 5.44 mm in the vertical direction. The shift register, decoder, transfer lead drivers, test circuits, and electronic I/O circuitry surround the smart pixel elements of the array. The total chip size is $7\text{-mm} \times 7\text{-mm}$.




Fig. 5. Reflectivity versus voltage for the MQW diodes on the four corners of the array, measured at ~ 851.5 nm and ~ 20 mW.

The chip is packaged on a custom aluminum mount with a four-layer flexible microstrip circuit providing the control signal and bias connections. Bypass capacitors are connected between the V_{dd} , $V_{P-ECL}(3\text{ V})$, $V_{modulator}$, and $V_{detector}$ power supplies and ground, and 50-ohm resistors are connected between the input control signals and ground. To interface with standard test equipment, we chose to connect the resistors to ground rather than to V_{P-ECL} .

All but two of the optical modulators and detectors generated photocurrent in response to incident light. The voltage dependence of the reflectivity and responsivity of the modulators and detectors can be measured by sweeping the voltage between the detector power supply and ground. Although there is no direct connection of the photodiodes to ground, the receiver circuitry completes the connection. First, a forward biased diode in series with the photodiodes, from the parasitic diodes between the p-diffusion and n-well of the p-type feedback FET, that provides a current path to V_{dd} . A nonlinear "resistor," from the static current versus voltage characteristics of the receivers and electrical differential amplifiers, completes the connection between V_{dd} and ground. The reflectivities in Fig. 5 are flat for voltages below 2.5 V, because the 0.7-V drop across the forward-biased diode and 1.8-V drop from V_{dd} to ground reduce the voltage that appears across the photodiode from that supplied to the circuit. Nonetheless, the data can be used to compare the uniformity of the detectors and to measure high and low state reflectivities.

The reflectivities for the detectors on the four corners of the array are shown in Fig. 5. The thickness of the antireflection coating was not optimum, limiting the contrast ratio to less than 2:1. The high and low state reflectivities at 6 and 11 V were 0.44 ± 0.03 and 0.25 ± 0.03 . The uniformity in reflectivities is much better than in our previous devices [10] because of better uniformity of the thickness of the stop etch layer. If the A/R coating is not perfect, Fabry-Perot resonances will be present, and variations in cavity length will vary the resonant frequency, which in turn varies the reflectivity at a fixed wavelength.

In Fig. 6, we show the normalized output from one 16×1 node from each of the sixteen 16×16 sections in the array, with two of the 16 inputs active at a data rate of 200 Mb/s. Input 0 had a pattern of "11 100010" and input 4 (0100)




Fig. 6. Normalized output bit patterns at 200 Mb/s from one 16×1 node from each of the different 16×16 switches (see Fig. 1). Input 0 had a pattern of "11 100010" and input 4 had a pattern of "01010011".

had a pattern of "01010011." In this figure, all devices had the correct output bit pattern. We also measured all 256, 16×1 nodes with four of 16 inputs active at 200 Mb/s. One of the 20 decoded control bits was stuck in a fixed state, preventing selection of eight of the inputs (two of which were measured) from the nodes in the bottom row of the array. Other than that, all devices had easily recognizable bit patterns. There was significant output amplitude variation from device to device, caused by defocus and positioning errors as the motorized stages moved the array across the fixed input and read beams. In particular, the defocus caused the amount of light coupled into the output fiber-based detector to be reduced. The data in Fig. 5 indicates this variation is not present in the devices, because the reflectivities of the modulators are fairly uniform and they are driven by voltages that should not vary in amplitude. Normalizing the data in Fig. 6 caused the "noise" to be magnified toward the lower right corner where the amplitude was reduced. The delay variation in these measurements of ~ 1 ns was also likely caused by positioning errors leading to variations in photocurrent. More detailed measurements on uniformity and delay variations and are given below.

In Fig. 7, we show a superposition of 16 eye diagrams at 400 Mb/s of a $16:1$ node with one input selected at a time. The photocurrent was monitored as the input spot was moved from detector to detector to ensure positioning errors did not contribute to eye closure. The eye depicts the combined jitter, skew, and pulse width distortion for an entire node and has sufficient opening for reliable operation. One could achieve a clean eye diagram at data rates up to ~ 470 Mb/s. The limiting factor is the modulator driver which was an inverter with $3\text{-}\mu\text{m}$ wide FET's. There was not additional space to make a larger driver.

Next, we looked at the dependence of the pulse width on the node bias voltage (V_{dd}) and on optical power. Pulse

TABLE I
CALCULATED CHIP DISSIPATION FOR VARIOUS SECTIONS OF THE CHIP

Element	Cap (pF)	Rate (Mb/s)	transitions per bit	diss./line (mW)	number	diss total (mW)
Control load	3	25	0.031	0.029	65	1.88
transfer bit	3	25	0.031	0.029	68	1.96
decoder bits	0.8	25	1	0.250	10	2.50
control information	3.7	25	1	1.156	68	78.63
input pad (dynamic)	2	25	1	0.625	23	14.38
idle channel	2.5	50	1	1.563	4	6.25
clock	3	25	2	1.875	2	3.75
rcvr-on (dynamic)	0.12	200	1	0.300	256	76.80
modulator(dynamic)	0.06	200	1	0.150	256	38.40
rest of node	0.25	200	1	0.625	256	160.00
rcvr-on (static)		200		2.500	256	640.00
rcvr-off (static)		200		0.060	3840	230.40
modulators (static)		200		0.400	256	102.40
input pads(static)		25		6.000	45	270.00
Total dissipation						1627.34

The calculations are figured assuming the entire chip had optical inputs and was operating. The dynamic dissipation is equal to $1/2NCV^2B_{eff}$, where C is the capacitance, V is the voltage swing (5 V), B_{eff} is the effective bit rate, which is equal to the actual bit rate (200 Mb/s for data, 50 Mb/s for idle channel, and for, 25 Mb/s control), times the number of transitions per bit (for the control loads, we assume that there is one pulse or two transitions every 65 bits, for RZ data there are two transitions per bit, for NRZ data there is one transition per bit), and N is the number of lines of the chip. The capacitance was estimated by summing the total gate area capacitance and twice the extracted values (to be conservative). Static power dissipations were simulated using SPICE. We do not know the exact dissipation of the disabled receivers that were not illuminated.




Fig. 7. Sixteen eye diagrams at 400 Mb/s superimposed, where each eye diagram is the optical output from a 16×1 switching node with one of its 16 optical inputs illuminated with pseudorandom data with a word length of 2^{23} .

width distortion places a limit on the maximum bit rate that can be achieved in the system application of these arrays. The measured pulse widths were different for the devices in the even and odd columns of the array, because two different feedback resistors were used. This was accomplished by varying the gate length of the p-type FET from 1–1.5 μm . The feedback resistor determines the optical input power or current that causes the output of the receiver to change from a low to high value (i.e., the receiver threshold). The circuit was designed with two different feedback resistor values in case the threshold ended up too high and we did not have enough optical power or in case the threshold ended up too low and the RC time constant was too long. This was unnecessary, because the threshold can be adjusted by varying V_{dd} . The dependence on V_{dd} occurs because the threshold of the receiver NAND gate is a function of V_{dd} , so the nominal gate to source voltage of the p-type FET changes as a function of V_{dd} , and thus




Fig. 8. Solid lines show pulse width versus optical input power at various values of V_{dd} from 5.0–6.0 V for a node with the smaller effective feedback resistance value near the upper left corner of the array. Solid circles indicate the actual data points. Power was assumed to be twice the photocurrent, which was monitored during the set of measurements. The dotted lines, with point labels +, o, and x, correspond to nodes near the other three corners of the array at $V_{dd} = 5.4$ V. The time resolution of the measurements was 198 ps.

its effective resistance changes. Higher values of V_{dd} should increase the gate to source voltage and thus lower the feedback resistance, thereby raising the effective optical power threshold of the gate. Thus for a given delay or pulse width, higher optical powers are needed for higher V_{dd} . This was confirmed experimentally.

We modulated the input lasers at 200 Mb/s with a pattern consisting of "0 000 100 011 110 111." This pattern gives a lone "1" (the fifth bit) and a lone "0" (the 13th bit) Looking

at the width of these bits gives a good indication of pulse width distortion. Fig. 8 shows the pulse width of a node near the upper left corner of the array with one particular input selected for the fifth bit in the pattern (the lone "1") for various values of V_{dd} versus optical input power. The time resolution of the data was 195 ps. One would expect that if the lone "1" bit had a longer pulse width, the lone "0" bit would have a shorter pulse width and that the average of the two would be 5 ns. Indeed, the average of the two pulse widths was 5 ns to within 310 ps. The three dotted lines in the figure show the same data for nodes near the other three corners of the array, measured over a smaller power range. The total variation of the four nodes for a given optical power and voltage was less than ± 400 ps. The variation can be caused by differences in the transistor characteristics across the array or by variations in V_{dd} or ground potential across the array. Because the static current is low, it is likely the former. It is unlikely that random variations exist across the array, because the transistor characteristics tend to vary in a smooth fashion for an established CMOS process.

The data in Fig. 8 also shows that the allowed variations in optical power is greater for higher V_{dd} values or higher thresholds, even when plotted on a logarithmic scale. This occurs because the nonlinearity in the feedback resistor "compresses" the voltage swing as a function of power. For example, at high powers, a 20% increase in optical power might only cause a 10% increase in voltage swing, whereas at lower powers, this same increase in optical power would cause a 20% increase in voltage swing. Because the delay and pulse widths are functions of the voltage swing, a reduction in voltage swing versus optical power translates into less variation in pulse width as a function of optical power.

The measured static power dissipation of the array ranges from 750 mW at $V_{dd} = 5$ V to 1.5 W at $V_{dd} = 6$ V. Roughly 70% of this is in the receivers and 30% is in the input differential amplifiers for the electrical control signals. The dynamic dissipation can be estimated by $1/2 CV^2 B_{eff}$ where B_{eff} is the effective bit rate of the particular signals. In Table I, we show the design data rates for the various parts of the circuits and the calculated dynamic dissipations and compare them to the static dissipations. The static dissipations were simulated using SPICE. The unselected receiver dissipation, assuming that input light is present, is equal to the product of the photocurrent and the difference between the detector voltage and V_{dd} , because the parasitic diodes of the feedback FET's provide a path for the photocurrent to V_{dd} . For an average photocurrent of 20 mA and a voltage difference of 3.0 V, the 3840 unselected receivers have a dissipation of 230 mW. The static dissipation from the 256 optical modulators, for a photocurrent of 50 mA with a voltage of 8.0 V, is 102 mW. From the table, one can see that the ~ 2.5 -mW static power dissipation of the receivers dominates. This does not mean that optical interconnections are not warranted, indeed, the use of optical interconnections greatly reduces dynamic dissipation by eliminating the need for long electrical traces on the chip (and large transistors to drive them) as well as large-dissipation-hungry electronic output drivers. It may be more optimum from a power dissipation point of view to have

more electronics per optical I/O. Others have also reached this conclusion [14].

Incidentally, during testing, only one or two nodes had optical inputs. Without optical inputs, the static dissipation of the unselected receivers and modulators and the dynamic dissipation of the switching nodes (including the receivers and modulators) are both approximately zero. For this reason, it is important to build systems to test the entire array concurrently to demonstrate that power dissipation will not be a problem.

Two cases of crosstalk were measured using two lasers with slightly different bit rates incident on the devices. In one case, the interfering signal was incident on detectors within the same node, and in the other case the interfering laser was incident on detectors in a neighboring node. With the pseudorandom optical data inputs at 200 Mb/s, no eye closure is observed in either case. However, by looking at the detected optical output on a spectrum analyzer, crosstalk, 45 dB below the signal, could be observed at 200 MHz for square wave inputs when the interfering signal was incident on a selected receiver of a neighboring node. No observable crosstalk was seen when the signal was incident on an unselected receiver, either in the same node or a neighboring node.

Voltage variations due to simultaneous switching currents through the parasitic inductances and resistances of the supply lines is the most likely reason for the observed crosstalk. The modulator driver supply lines are more likely to contribute crosstalk than as the receiver supply lines, because the crosstalk was independent of the selected receiver on the neighboring node. If we extrapolate the crosstalk value measured and assume each of the 16 nodes in a row on a common bias lead will contribute the measured amount, the overall signal to noise ratio should be $-45 \text{ dB} + 10 \log(16) = -33 \text{ dB}$. This would cause only a 0.20-dB power penalty for an input noise limited receiver with an incident signal that comes from the node output.

In conclusion, we have presented the first high-speed optoelectronic VLSI switching chip using flip-chip bonded MQW modulators on VLSI electronics. It implements 16 complete 16×16 switches operating above 400 Mb/s per channel. The chip has more than 140-K transistors and more than 4-K optical I/O. The yield and uniformity of the chip far exceeds that made using monolithic GaAs-based technologies and is acceptable for use in systems experiments. Indeed, a recent demonstration system operating from 155–208 Mb/s has been made using this chip [11]. The silicon circuitry is made using mature technology, and state of the art silicon technologies should enable chips to be designed with per channel speeds over 1 Gb/s and throughputs of 1 Tb/s.

ACKNOWLEDGMENT

The authors would like to thank T. K. Woodward, A. V. Krishnamoorthy, and T. J. Cloonan for helpful discussions on circuit design.

REFERENCES


- [1] I. Underwood, D. G. Vass, A. O'Hara, D. C. Burns, P. W. McOwan, and J. Gourlay, "Improving the performance of liquid crystal spatial light modulators: Issues and achievements," *Appl. Opt.*, vol. 33, pp. 2768–2774, 1994.

- [2] D. J. McKnight, K. M. Johnson, and R. A. Serati, "256 × 256 liquid crystal on silicon spatial light modulator," *Appl. Opt.*, vol. 33, pp. 2775–2784, 1994.
- [3] L. A. D'Asaro, L. M. F. Chirovsky, E. J. Laskowski, S. S. Pei, T. K. Woodward, A. L. Lentine, R. E. Leibenguth, M. W. Focht, J. M. Freund, G. D. Guth, and L. E. Smith, "Batch fabrication and operation of GaAs-AlGaAs field-effect transistor self electro-optic effect device (FETSEED) smart pixel arrays," *IEEE J. Quantum Electron.*, vol. 29, pp. 670–677, 1993.
- [4] J. Wieland, H. Melchoir, M. Q. Kearly, C. Morris, A. J. Mosley, M. G. Goodwin, and R. C. Goodfellow, "Optical receiver array in silicon bipolar technology with self-aligned low parasitic III/V detectors for DC-1 Gb/s parallel links," *Electron. Lett.*, vol. 27, p. 2211, 1991.
- [5] C. Camperi-Ginestet, M. Hargis, N. Jokerst, and M. Allen, "Alignable epitaxial lift-off of GaAs materials with selective deposition using polyimide diaphragms," *IEEE Photon. Technol. Lett.*, vol. 3, p. 1123, 1991.
- [6] K. W. Goossen, J. A. Walker, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller, "GaAs MQW modulators integrated with silicon CMOS," *IEEE Photon. Technol. Lett.*, vol. 7, pp. 360–362, 1995.
- [7] T. L. Worchesky, R. J. Ritter, R. Martin, and B. Lane, "Large arrays of spatial light modulators hybridized to silicon integrated circuits," *Appl. Opt.*, vol. 35, pp. 1180–1186, 1996.
- [8] T. K. Woodward, A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Jan, L. A. D'Asaro, L. M. F. Chirovsky, S. P. Hui, B. Tseng, D. Kossives, D. Dahringer, and R. E. Leibenguth, "1 Gb/s two-beam transimpedance smart pixels optical receivers made from hybrid GaAs MQW modulators bonded to 0.8 μ m silicon CMOS," *IEEE Photon. Technol. Lett.*, vol. 7, pp. 422–424, 1996.
- [9] A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, T. K. Woodward, J. E. Ford, G. F. Aplin, L. A. D'Asaro, S. P. Hui, B. Tseng, R. E. Leibenguth, D. Kossives, D. Dahringer, L. M. F. Chirovsky, and D. A. B. Miller, "3-D integration of MQW SEED modulators over active sub-micron CMOS circuits: 375 Mb/s transimpedance receiver-transmitter circuit," *IEEE Photon. Technol. Lett.*, vol. 7, pp. 1288–1290, 1995.
- [10] A. L. Lentine, K. W. Goossen, J. A. Walker, L. M. F. Chirovsky, L. A. D'Asaro, S. P. Hui, B. Tseng, R. E. Leibenguth, D. P. Kossives, D. W. Dahringer, D. D. Bacon, T. K. Woodward, and D. A. B. Miller, "Arrays of optoelectronic switching nodes comprised of flip-chip bonded MQW modulators and detectors on silicon CMOS circuitry," *IEEE Photon. Technol. Lett.*, vol. 7, 1996.
- [11] A. L. Lentine, D. J. Reiley, R. A. Novotny, R. L. Morrison, J. M. Sasian, M. G. Beckman, D. B. Buchholz, S. J. Hinterlong, T. J. Cloonan, G. W. Richards, and F. B. McCormick, "Optoelectronic ATM switch employing hybrid silicon CMOS/GaAs FET-SEED's," in *Optical Interconnects in Broadband Switching Architectures*, T. J. Cloonan, Ed., vol. Proc. SPIE-2692, pp. 100–108, 1996.
- [12] M. Bazes, "Two novel fully differential complementary self-biased CMOS differential amplifiers," *IEEE J. Solid-State Circuits*, vol. 26, pp. 165–168, 1991.
- [13] J. M. Kuo, Y. C. Wang, S. P. Hui, B. J. Tseng, L. M. F. Chirovsky, K. W. Goossen, R. E. Leibenguth, G. Livescu, W. Y. Jan, J. E. Cunningham, J. A. Walker, L. A. D'Asaro, A. Ron, and W. E. Mayo, "Large area substrate removal using InGaP as the etch stop for a GaAs detector/modulator array flip-chip solder bonded to silicon CMOS," *IEEE Photon. Technol. Lett.*, to be published.
- [14] A. V. Krishnamoorthy, P. J. Marshand, F. E. Kiemelev, and S. C. Esener, "Grain size considerations for optoelectronic multistage interconnection networks," *Appl. Opt.*, vol. 31, pp. 5480–5506, 1992.

Keith W. Goossen (S'87–M'88) received the B.S. degree in electrical engineering from the University of California, Santa Barbara, in 1983, and the Ph.D. degree in electrical engineering from Princeton University, Princeton, NJ, in 1988. In his doctoral thesis he invented the first photovoltaic 10-micron detector based on photoemission from quantum wells.

In 1988, he joined AT&T Bell Laboratories, Holmdel, NJ. While there, his work has been primarily on surface-normal optical modulators and their integration with electronics for optical interconnects. This includes the first demonstration of high-speed, VLSI-density optical transceivers utilizing silicon CMOS, and GaAs-based modulators via flip-chip bonding. He has also investigated integration of quantum-well modulators on silicon, via direct epitaxy, and produced extensive research on quantum-well modulator technology, including the discovery of strong excitons in shallow quantum wells, and the development of stacked diode modulators and wafer-scale Fabry–Perot modulators. He has also invented the optical principle of the first broadband micromechanical modulator for local access fiber networks. In addition, he has developed modeling techniques for electrical propagation on normal and superconducting microstrip transmission lines.

J. A. Walker, photograph and biography not available at the time of publication.



Leo M. F. Chirovsky (M'95) was born in Munich, Germany, in 1948. He received the B.S. degree in physics from Cooper Union, New York, NY, and the M.A. and Ph.D. degrees from Columbia University, New York, NY, in 1970, 1975, and 1979, respectively.

In 1979–1980, he was an NSF Postdoctoral Fellow at Columbia University. In 1980, he joined AT&T Bell Laboratories, Murray Hill, NJ, where he worked on magnetic bubble memories, IC content-addressable memories, and parametric amplifiers.

Since 1986, he has been working on quantum-well optoelectronic devices and circuits. He holds several patents and has co-authored numerous papers in this field.

Dr. Chirovsky is a member of the American Physical Society and the Optical Society of America.



L. Arthur D'Asaro (S'49–M'55–SM'66–LS'93) received the B.S. and M.S. degrees in physics from Northwestern University, Evanston, IL, in 1949 and 1950, respectively, and the Ph.D. degree in engineering physics from Cornell University, Ithaca, NY, in 1955.

Since that time, he has been with AT&T Bell Laboratories where he has been mainly concerned with the exploratory development of semiconductor devices. Currently, he is participating in the development of optoelectronic switching device arrays.

Dr. D'Asaro is a member of the American Physical Society, Sigma Xi, and Phi Beta Kappa.

S. P. Hui, photograph and biography not available at the time of publication.

Anthony L. Lentine (M'83), for photograph and biography, see this issue, p. 2.

B. J. Tseng, photograph and biography not available at the time of publication.

R. E. Leibenguth, photograph and biography not available at the time of publication.

J. E. Cunningham, photograph and biography not available at the time of publication.

W. Y. Jan, photograph and biography not available at the time of publication.

Jen-Ming Kuo (S'85-M'87), photograph and biography not available at the time of publication.

D. W. Dahringer, photograph and biography not available at the time of publication.

D. P. Kossives, photograph and biography not available at the time of publication.

D. D. Bacon, photograph and biography not available at the time of publication.




Gabriela Livescu (M'95) was born in Timisoara, Romania. She received the B.S. degree in physics from the University of Bucharest, Romania, in 1974, and the Ph.D. degree in physics from the Technion, Israel, in 1984.

She was a Research Associate at the City College of New York between 1984-1986, after which she joined AT&T Bell Laboratories (now Lucent Technologies), where she is now a Member of the Technical Staff. Her research at Bell Laboratories covered various aspects of quantum-well materials

and devices, ranging from very fundamental to systems and applications. In basic research, she is best known for her work on free-carrier and many-body effects on the optical properties of modulation-doped structures, as well as for her time-resolved studies on charge transport and tunneling in quantum wells, and on enhanced photoluminescence from DX centers in AlGaAs heterostructures. During the last several years, her work has shifted toward devices and their applications, most recently focusing on smart pixels, with emphasis on high-speed operation and optical characterization of large arrays for optical interconnections. At present, she is involved in work on optical systems for laser communications.

Dr. Livescu is a Member of the Optical Society of America and the American Physical Society.

R. L. Morrison, photograph and biography not available at the time of publication.



Robert A. Novotny (M'88) received the B.S.E.E. and M.S.E.E. degrees from the Illinois Institute of Technology, Chicago, IL, in 1984 and 1986, respectively.

In 1978, he joined AT&T Bell Laboratories, Naperville, IL, and worked in the development of voice and data packet switching systems. In 1989, he joined the Broadband and Photonic Switching Department, where he investigated the use of SEED modulators and VCSEL's in two-dimensional optical data links. His recent work is in the design

and modeling of optoelectronic integrated circuits for use in large ATM switching systems.

D. B. Buchholz, photograph and biography not available at the time of publication.

Analysis of Parasitic Front-end Capacitance and Thermal Resistance in Hybrid Flip-chip-bonded GaAs SEED/Si CMOS Receivers

R.A. Novotny, A.L. Lentine, D.B. Buchholz, A.V. Krishnamoorthy*,

AT&T Bell Laboratories, 2000 N. Naperville Rd., Naperville, IL 60566;

*Holmdel, NJ

(708) 713-5419, FAX (708) 713-7951

Smart pixels^[1] consisting of photodetectors, electronic circuitry, and E/O converters utilizing free-space optical interconnections show promise to relieve the interconnection bottleneck in computing and switching systems.^[2] To reduce the propagation delay through a smart pixel, the receiver requires a fast response, hence it is essential to reduce the front end capacitance (C_{in}). C_{in} has three main components: the photodiode active area, the amplifier input, and the stray interconnect capacitance (C_s). The FET-SEED technology minimizes C_s through the monolithic integration of photodetectors, modulators and electronic circuitry.^{[3][4]} However, current system demonstrations using FET-SEEDs have been limited to using medium scale

integration (MSI) smart pixel arrays. Hybrid integration of VLSI Si CMOS electronic circuitry with photodetectors, modulators, or emitters is an attractive approach in obtaining VLSI smart pixels in the near term.

One method of attaching III-V devices to Si CMOS is through the use of a flip-chip solder bump process and back illuminating the photodiode.^[5] A recent technique has been devised where GaAs SEED detectors/modulators are first flip-chip-bonded onto Si CMOS, and then the GaAs substrate is etched away allowing operation at 850nm.^[6] A question to be answered is what stray input capacitance results from this process.




Figure 1: Cross-sectional view depicting flip-chip hybrid, along with the equivalent circuit. (not to scale)

This paper investigates C_s as a function of solder bump height and diameter, using the current process' design rules. Figure 1 depicts the cross-sectional view of the flip-chip hybrid model along with the equivalent circuit. The current design rules dictate that the pads be equally sized squares spaced one pad width apart. Circuits with $15\mu\text{m}$ pads have recently been demonstrated.^[11] The SEED chip has a fixed $\sim 2\mu\text{m}$ overhang beyond the pad size, and the photodiode active area is slightly larger than one of the pads.

The total front end capacitance (C_{in}) was first estimated by taking the sum of all the contributing elements: $C_{in} = C_{amp} + C_{diode} + C_s$. Where $C_s = C_{trace} + C_{pad} + C_{chip} + C_{bump}$. The formulas used to approximate each element are listed in Table 1. Figure 2 plots the estimated C_{in} (less the fixed amplifier contribution) vs. pad size. Our results indicate that the pad was the dominate contributor to C_s . Solder bump heights from $5\text{--}20\mu\text{m}$ were found to induce little change on C_s .

To check the accuracy of the approximations, a 3-D Laplace/Poisson solver was used to calculate the total input capacitance vs. pads size for a SEED bumped to the first layer metal of a Si wafer. The results are shown in Figure 3, and had less than 2% error in symmetry preservation of the resulting Maxwell capacitance matrix. The small shaded region indicates solder bump heights ranging from $5\text{--}20\mu\text{m}$. The results agree reasonably well with




Figure 2: Plot of estimated input capacitance as a function of bond pad size.

the estimated values (reshown as a dotted line in Figure 3) which appear to underestimate the fringing components of the structure.

To verify the above simulations, CMOS ring oscillators have been designed with and without solder bumped SEED loads. Test results will be discussed.

The effect of thermal conduction from the SEED to Si substrate was also examined. The output contrast of a SEED modulator diminishes with change in temperature due to the shift of the exciton ($0.28\text{nm}/^\circ\text{C}$). The amount of heat generated in the SEED is dependent on the impinging optical power (P_{in}), and its state of absorption. Light not reflected is absorbed

ELEMENT	APPROXIMATION	DESCRIPTION
C_{amp}	25fF	Assumed amplifier input capacitance
C_{trace}	1.2fF	Interconnect to amp is a fixed $2 \times 5\mu\text{m}$ trace ^[8]
C_{diode}	$(K_s d^*(d+2))(115\text{aF}/\mu\text{m}^2)$	SEED active area ^[7] (Fringing factor $K_s \approx 0.6(1/d+1)$)
C_{pad}	$d^2(0.031 \text{ fF}/\mu\text{m}^2) + 4(d)(0.044\text{fF}/\mu\text{m})$	Metal-1 to substrate + fringing ^[8]
C_{bump}	$63.5\epsilon_r \text{ aF}$	Capacitance between two spheres radius= $r(\mu\text{m})$ ^[9]
C_{chip}	$(K_c)(\epsilon_r)(\epsilon_o)(d(d+4)/h)$	Conductor over a ground plane (GaAs chip over Si) ^[10] $K_c = \text{fringing factor} \approx (1.1h/d + 1)$ for $h/w < 2$

TABLE 1: Formulas used for the approximation of C_{in} .




Figure 3: Plot of simulated input capacitance as a function of bond pad size. Shaded region indicates solder bump heights from 5-20 μ m. For comparison, the estimated capacitance is shown as a dotted line.

as a photocurrent, which generates heat. Assuming a modulator biased at 6V, with $P_{in} = 500\mu W$, and a high/low state differential responsivity of 0.2/0.6A/W, results in $P = 1.2mW$ differential in heat dissipation between the two states. Figure 4 shows the thermal network used to model heat conduction. For a 15 μ m square pad and 10 μ m bump height, the following values were estimated:^[12] $R_{GaAs} = 19.2k$, $R_{bump} = 1.23k$, $R_{SiO2} = 0.44k$, $R_{total} = (R_{bump} + R_{SiO2}) \parallel (R_{bump} + R_{SiO2} + R_{GaAs}) = 1.67k \parallel 20.9k = 1.54k$. The change in temperature of SEED due to photo current would be: $\Delta T = (\Delta P)(R_{total}) = (1.2mW)(1.54k) = 1.85^\circ C$. This would result in a negligible drop (<.2dB) in output contrast. Thus, the hybrid smart pixel




Figure 4: Diagram depicting the thermal network of SEED/Si hybrid.

technology examined here has both acceptable thermal and electrical performance for the current design rules.

Acknowledgments

This work was partially sponsored by ARPA under Air Force Rome Laboratories contract No. F30602-93-C-0166

References

- [1] See "Summer Topical Meeting Digest on Smart Pixels," IEEE Cat No. 94TH0606-4
- [2] J.W. Goodman, "Switching in an Optical Interconnect Environment," 1989 OSA Proceedings on Photonic Switching, Vol 3., March 1989]
- [3] D.A.B. Miller, M.D. Feuer, Y.T. Chang, A.S.C. Shunk, J.E. Henry, D.J. Burrow, D.S. Burrows, D.S. Chemla, "Field-Effect Transistor Self-Electrooptic Effect Device: Integrated Photodiode, Quantum Well Modulator and Transistor," IEEE PTL, V1, N3, March 1989.
- [4] L.A. D'Asaro, L.M.F. Chirovsky, E.J. Laskowski, S-S. Pei, R.E. Leibenguth, T.K. Woodward, M. Focht, A.L. Lentine, M.T. Asom, G. Guth, R.F. Kopf, J.M. Kuo, S.J. Pearton, G.J. Przybylek, F. Ren, L.E. Smith, "Batch Fabrication and Structure of Integrated GaAs-AlGaAs Field-Effect Transistor-Self Electro-optic Effect Devices (FET-SEEDs)," IEEE EDL, V13, N10, October 1993.
- [5] R.S. Sussmann, R.M. Ash, A.J. Moseley, R.C. Goodfellow, "Ultra-low Capacitance Flip-chip-bonded GaInAs PIN Photodetector for Long Wavelength High-data-rate Fiber Optic Systems," Elec Lett., V21, N14, July 1985]
- [6] Goossen KW; GaAs 850-nm Modulators Solder Bonded to Silicon; IEEE PHOTON V5, 7, July 1993
- [7] A.L. Lentine, L.M.F. Chirovsky, L.A. D'Asaro, C.W. Tu, D.A.B. Miller, "Energy Scaling and Subnanosecond Switching of Symmetric Self-Electrooptic Effect Devices," IEEE PTL, V1, N6, June, 1989.
- [8] MOSIS wafer acceptance specifications for HP CMOS26B 0.8 micron CMOS bulk wafers.
- [9] C.W. Walker, "Capacitance, Inductance, and Crosstalk Analysis," Artech House, Inc, ISBN: 0-89006-392-3, 1990 (derived from formula set on p83).
- [10] ibid, derived from formula set in section 2.2.6.
- [11] K.W. Goossen, A.L. Lentine, J.A. Walker, L.A. D'Asaro, S.P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. Dahringer, L.M.F. Chirovsky, D.A.B. Miller, "4x4 Array of GaAs Hybrid-on-Si Optoelectronic Switching Nodes Operating at 250Mbps," Paper PD2.2, Leos Annual Meeting, November 2, 1994
- [12] R.R. Tummala, E.J. Rymaszewski, "Microelectronics Packaging Handbook," Van Nostrand Reinhold, Ch 4, ISBN:0-442-20578-3, 1989.

THERMAL CONSIDERATIONS IN THE DESIGN OF OPTO-ELECTRONIC DEVICE MOUNTS

D.B. Buchholz, A.L. Lentine, R.A. Novotny
AT&T Bell Laboratories, 2000 N. Naperville Rd., Naperville, IL 60566
(708) 979-2809, FAX (708) 713-7951

For some optoelectronic devices, the optical characteristics of the device are a strong function of the device temperature. An example of this is the self-electro-optic effect device (SEED).^[1] SEEDs make use of the shift in wavelength of the exciton absorption maxima that occurs as a function of a changing electrical field across multiple quantum well material.^[2] For a typical device, the absorption maxima must be shifted by 3 to 5nm to obtain good contrast between the absorptive and reflective states. A change in the device temperature, however, can also change the location of the exciton absorption maxima.^[3] For GaAs/AlGaAs devices the absorption maxima shifts approximately $0.28\text{nm}/^\circ\text{C}$.^[4] Thus, it becomes necessary to carefully design the SEED mount to minimize the temperature gradient and, therefore, this temperature induced shift across the chip.

In this paper we have used finite element analysis (FEA) to model mounts for a 16×16 array of FET-SEED switching nodes. By careful mount design, the calculated temperature spread could be held to 1°C even when the power density was $40\text{W}/\text{cm}^2$ over the 0.15cm^2 active chip area; a 6W chip. We have also used the temperature dependence of the exciton absorption maxima to map the temperature of an existing 4×4 array of FET-SEED switching nodes^[5], operating at a power density of $49\text{W}/\text{cm}^2$ and found the results in good agreement with those obtained by FEA.

Two fundamental tasks exist in the thermal management of a chip; not only is it necessary to prevent the entire chip from heating to a point where thermal effects degrade the overall performance but it is also necessary to maintain the temperature of multiple locations on the chip to nearly the same temperature. The exact amount of temperature variation that can be tolerated (ΔT) will depend on the application but in general it will fall into the range

of 1 to 4°C . It is not necessary to hold the entire chip to this temperature range but only specific devices on the chip, such as all the optical output modulators. The temperature variation of interest is the difference between the hottest and coolest of these devices, as shown in Figure 1, not the overall temperature spread or the temperature differential of a single node. The temperature profile depicted




Figure 1
Schematic 1D Temperature Profile
SEED Node Array

in Figure 1 is representative of that which one would expect in a regular array of nodes. The chief cause of temperature variation between equivalent devices on the chip is the spreading of heat to the predominately passive regions around the periphery of the chip where the electrical I/O bond pads are located.

The design of a SEED mount must, therefore, contain an analysis of the temperature variation between the equivalent critical optoelectronic components on the chip. The testing of the mounted chip must contain an analysis to verify the temperature variation between these components. The first of these two tasks can be accomplished by finite element analysis and will be discussed later. The second of these two tasks can be accomplished using the same physical phenomena that makes attention to temperature variation necessary, namely the shifting of the exciton peak location. The experimental setup used in this study is shown in Figure 2. The chip, a 4×4 array of $210\mu\text{m} \times 210\mu\text{m}$ FET-SEED switching nodes^[5], was held at

a constant temperature with a thermal electric cooler (TEC). A laser light source, $\lambda=850\text{nm}$, and




Figure 2
Experimental Setup

lens system was used to illuminate a single SEED modulator. The reflected light was focused on a photodetector. A bias voltage ramp was applied to the modulator and the photocurrent of the reflected light measured as a function of the applied bias. The minimum photocurrent corresponds to the absorption maxima. Next, the temperature of the entire chip was changed by adjusting the TEC. The change in bias voltage at which the absorption maxima occurred, for the same modulator, was noted. In this manner a ΔV versus ΔT curve was constructed. Finally an X-Y stepper stage was used to move the laser illumination to different modulators and the bias voltage at which the absorption maxima occurred noted as a function of chip location. A finite element analysis (FEA) program was used to model the temperature profile of the measured device. The device mount is depicted in Figure 3. The thermistor used to control the TEC, was




Figure 3
Thermal Model for FEA

not part of the thermal model but is included to further define the experimental setup. The thermistor was held at 20°C . The chip was powered to 350mW and the power was assumed to be distributed evenly over the $840\mu\text{m} \times 840\mu\text{m}$ active area of

the array ($49\text{W}/\text{cm}^2$). The overall chip size was approximately 2.8mm square and the ceramic was 7.5mm square. Figure 4a is a plot of the measured temperature for four different modulators on the chip. Figure 4b is a plot of the isotherms on the




Figure 4a
Experimentally Determined
Temperature Profile
4x4 SEED Array




Figure 4b
Finite Element Analysis
Temperature Profile
4x4 SEED Array

active area of the chip surface as predicted by the FEA program. The thickness of the thermal epoxy used to affix the device to the ceramic was an estimate and could easily account for the minor differences between the model and the experimental measurements.

The chip just examined had a temperature variation across the active area on the order of 1.5°C , which for most applications would be acceptable. As mentioned, the chip was a 4×4 array of SEED nodes. For a chip of this size, even the center nodes are close to the non-active boarder regions and the effect of thermal spreading is therefore minimized. The temperature distribution for a 16×16 array of $240\mu\text{m} \times 240\mu\text{m}$ FET-SEED switching nodes, operating at 6.3W ($40\text{W}/\text{cm}^2$ over a 0.15cm^2 active area) mounted the same as the 4×4 array, is shown in Figure 5. Not only does the overall temperature




Figure 5
Finite Element Analysis Temperature Profile
16 x 16 Node SEED Array

of the chip rise but the temperature variation increases to 5°C . Again, the chief cause of the temperature variation is heat spreading into the inac-

tive boarder regions of the chip; the active chip area is 3.84mm x 3.84mm and the overall chip size is 4.32mm x 4.32mm. A mount for the 16x16 node array was designed that would counter the effects of the heat spreading, as shown in Figure 6. An




Figure 6
Heat Confining Chip Mount

opening is cut in the ceramic used for the device interconnect and the device is mounted on a molybdenum pedestal that tapers down as it approaches the TEC. The taper in the pedestal counteracts the effect of the inactive chip periphery and the temperature spread is reduced to approximately 1°C, as shown in Figure 7. The decrease in




Figure 7
Finite Element Analysis Temperature Profile
16 x 16 Node SEED Array
With Heat Confining Chip Mount

the overall chip temperature is due to the use of solder for the two bonds in the thermal path instead of organic adhesives. The use of low temperature solders is required because of temperature limitations imposed by the TEC.

In conclusion, we have mapped the temperature across the active area of a SEED array by calibrating the exciton peak shift, of the GaAs-AlGaAs quantum well modulators, as function of temperature. We have compared the measured temperature profile to one modeled by finite element analysis and found them in good agreement. Using the finite element analysis program, we have shown that with proper design of the chip mount, areas of equivalent optoelectronic functions can be held at a

uniform temperature, $\Delta T \sim 1^\circ\text{C}$, even for large high power density devices.

This work was partially sponsored by ARPA under Air Force Rome Labs contract number F30602-93-C-0166.

References

- [1] D. A. B. Miller, D. S. Chemla, T. C. Daman, A. C. Gossard, W. Weigmann, T. H. Wood, C. A. Burrus, "Novel hybrid optically bistable switch: The quantum well self-electro-optic effect device," *Appl. Phys. Lett.* **45**, pp. 13-15 (1984)
- [2] D. A. B. Miller, D. S. Chemla, S. Schmitt-Rink, "Electric Field Dependence of Properties of Semiconductor Quantum Wells: Physics and Applications", in *Optical Nonlinearities and Instabilities in Semiconductors*, H. Haug Ed., Academic Press, Inc., 1988
- [3] S. M. Sze, *Physics of Semi Conductor Devices*, 2nd Ed., John Wiley and Sons, 1981 pp.15
- [4] R. A. Novotney, "Parallel optical data links using VCSELs," *Proceedings SPIE*, vol. **2147**, *Vertical-Cavity Surface-Emitting Laser Arrays*, pp. 140-149 (1994)
- [5] D.B.A. Miller, M.D. Feuer, Y.T. Chang, A.S.C. Shunk, J.E. Henery, D.J. Burrow, D.S. Barrow, D.S. Chemla, "Field-Effect Transistor Self-Electrooptic Effect Device: Integrated Photodiode, Quantum Well Modulator and Transistor," *IEEE PTL*, **V1**, N3, March 1989

Arrays of Optoelectronic Switching Nodes Comprised of Flip-Chip-Bonded MQW Modulators and Detectors on Silicon CMOS Circuitry

A. L. Lentine, *Member, IEEE*, K. W. Goossen, *Member, IEEE*, J. A. Walker, L. M. F. Chirovsky, L. A. D'Asaro, *Life Senior Member, IEEE*, S. P. Hui, B. T. Tseng, R. E. Leibenguth, D. P. Kossives, D. W. Dahringer, D. D. Bacon, T. K. Woodward, *Member, IEEE*, and D. A. B. Miller, *Fellow, IEEE*

Abstract— We describe 8×8 arrays of smart pixels, designed and fabricated using MQW modulators and detectors flip-chip-solder-bonded to silicon CMOS circuits. The individual circuits implement 2 input, 1 output embedded control switching nodes. Four arrays from two different designs were fabricated and tested. For the array with the highest yield, 60 of 64 nodes functioned correctly at low speeds and were tested up to 250 Mb/s without re-adjusting individual bias voltages with the maximum speed of an individual node of 375 Mb/s. For the second-generation array, the center 4×8 section of the array was tested at data rates beyond 700 Mb/s with individual nodes having short term bit error rates below 10^{-11} .

ONE APPROACH to improving the performance of large processing or telecommunications switching systems is to interconnect integrated circuits using optics. Smart pixels, with integrated optical detectors, modulators, and electronic logic, could potentially be used in these systems. The FET-SEED, consisting of the monolithic integration of multiple quantum well (MQW) optical modulators and detectors with GaAs field effect transistors, is one design platform for these smart pixels [1], [2]. Another potential design platform uses the hybrid integration of MQW modulators and detectors with commercial electronic circuits [3]–[6]. This latter approach allows one to design circuits with greater complexity and circuit yield, because it uses available established VLSI processes.

We describe 8×8 arrays of smart pixels, designed and fabricated using MQW modulators and detectors flip-chip-solder-bonded to silicon CMOS circuits. The modulators were designed for 850 nm operation and the substrate was removed to avoid excess absorption in the substrate [5]. The individual circuits implement 2 input–1 output embedded control switching nodes.

The CMOS circuit shown in Fig. 1 is functionally similar to switching nodes previously made using the monolithic FET-

SEED technology [7]. Each node contains a single optical receiver. The first-generation arrays have two different receiver designs in alternating columns of the array, one with and one without voltage clamps on the receiver inputs. The second-generation arrays had four different transimpedance receiver designs with feedback elements shown in Fig. 1(b). The transimpedance receivers operated at lower optical powers and higher data rates. In both designs, the electrical output from a given receiver is connected to the data input of a first 2×1 multiplexer physically located within the same node as the receiver and a second 2×1 multiplexer physically located in a second node next to the first node. Each multiplexer has a pair of complementary electrical inputs, known as the control bit, that determines which input is regenerated as the optical output. In each node, a control memory (set-reset latch) stores this control bit. In the embedded control architecture, the format of the input optical signals consists of the control bits followed in time by the data bits. An electrical control load signal, common to all the nodes within the array, is held high to enable the writing of the control memories with the control bits. Once the control bits are loaded, the control load signal is held low to disable the writing of the memory, and the correct input data bits are regenerated at the output based upon the state of the memory. The output modulators are driven by an electrical inverter following the multiplexers. Other than the receivers, the circuit schematics were the same for the two designs, except that the FET's were wider in the second-generation design to provide increased current for operation at higher data rates.

The first- and second-generation circuits were designed using $1.2 \mu\text{m}$ and $1.0 \mu\text{m}$ CMOS. The center to center spacing of the nodes was $135 \mu\text{m} \times 120 \mu\text{m}$, so each node occupies an area equal to $\sim 28\%$ of that of the comparable monolithic circuit [7]. The bump-bond pad sizes and optical window sizes were $15 \mu\text{m} \times 15 \mu\text{m}$ with a minimum space between two bump-bonds of $15 \mu\text{m}$. Transistors were located greater than $2 \mu\text{m}$ from the bond pads, but recently circuits have been made with FET's directly under the pads [8]. The detector and modulator window spacings were $60 \mu\text{m}$ and the spacings between detectors and modulators within a node were $75 \mu\text{m}$.

The MQW modulators were made using processes similar to the monolithic FET-SEED [2]. The layer structure consisted of 95 periods of 9 nm GaAs quantum wells with 3.5 nm $\text{Al}_{0.3}\text{Ga}_{0.7}\text{As}$ barriers. Additional steps to the process included

Manuscript received August 9, 1995; revised October 25, 1995. This work was supported in part by ARPA under Air Force Rome Laboratories contract number F 30602-93-C-0166.

A. L. Lentine is with AT&T Bell Laboratories, 2000 N. Naperville Rd., Naperville, IL 60566 USA.

K. W. Goossen, J. A. Walker, T. K. Woodward, and D. A. B. Miller are with AT&T Bell Laboratories, 101 Crawfords Corner Rd., Holmdel, NJ 07733 USA.

L. M. F. Chirovsky, L. A. D'Asaro, S. P. Hui, B. T. Tseng, D. P. Kossives, D. W. Dahringer, and D. D. Bacon are with AT&T Bell Laboratories, 600 Mountain Ave., Murray Hill, NJ 07974 USA.

R. E. Leibenguth is with AT&T Bell Laboratories, 9999 Hamilton Blvd., Breinigsville, PA 18031 USA.

Publisher Item Identifier S 1041-1135(96)00944-5.

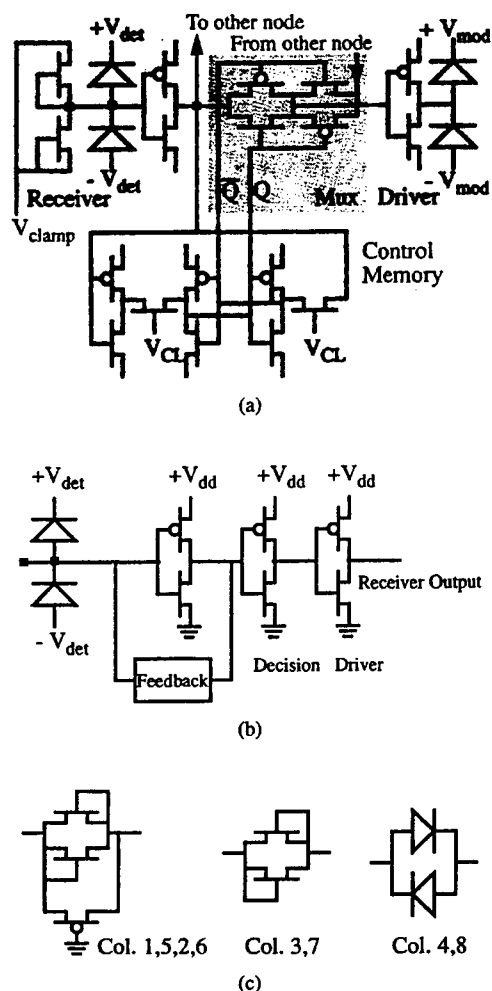


Fig. 1. (a) Schematic diagram of embedded control 2×1 nodes of the first-generation design. T 's are connected, and crosses are open unless indicated. n -fets are connected to GND, and p -fets are connected to V_{dd} . Clamping transistors are present on alternate columns. (b) Receiver section for second-generation design (c) feedback elements for receivers in (b) for the different columns in the array.

a deep mesa etch between diodes and the deposition of solder on the pads. After receiving the fabricated CMOS chips, additional metal layers (Ti, Ni, Au) and solder were deposited on the solder bump pads. After the GaAs chip was bonded to it, the GaAs substrate was removed, and the device was packaged and antireflection coated. A more detailed description of the process is given in [5]. A cross sectional schematic and photograph of a section of a bonded chip is shown in Fig. 2. The total height of the front of the modulator from the surface of the CMOS chip is $\sim 10 \mu\text{m}$.

Three arrays of the first design and one of the second were fabricated and tested. During substrate removal, the etchant attacked some of the modulators on the end columns of the array. For one array (of the first design), all but 1 detector and 3 modulators out of 256 quantum well diodes were operational after substrate removal.

Reflectivity and responsivity were measured for this array as a function of voltage for the bonded MQW diodes. The peak responsivities varied between ~ 0.4 and $\sim 0.45 \text{ A/W}$, the high and low state reflectivities varied from ~ 0.3 – ~ 0.4 and from ~ 0.06 – ~ 0.15 respectively at a fixed wavelength of ~ 850





Fig. 2. Cross sectional schematic (a) and photograph (b) of a section of the array. Rectangles are the individual MQW diode mesas, which measure $\sim 20 \mu\text{m} \times 50 \mu\text{m}$.

nm, and the contrast ratios varied from $\sim 1.8:1$ to $\sim 2.9:1$ for a 5 V swing.

High-speed testing was done on the arrays by current modulating the two input laser diodes with complementary sets of nonreturn to zero (NRZ) data from a digital word generator and supplying these optical inputs to one receiver at a time in the array. The center 4×5 , 6×6 , and 6×8 sections of the first design operated above 250 Mb/s and individual nodes were tested to 375 Mb/s. Fig. 3 shows the one of the outputs from each node from the center 4×8 section of the second-generation array at 700 Mb/s, with the control set so that each 2×1 node selected its own receiver. However, in all four arrays, we observed the same performance when either input of a given node was selected. At 700 Mb/s, the feedback resistors in columns 5 and 6 had too high of an impedance to affect the circuit, so the response was similar to column 3 that has FET clamps only. The nodes in column 4 with diode clamps required very asymmetric input powers, and the cause of this is unknown.

Fig. 4 shows one of the optical outputs from a node where the optical inputs to the receiver were modulated with a 10^{-23} pseudorandom pattern at 700 Mb/s. There is noticeable pulse-pattern dependency as evidenced by the separation of traces on the falling edges. This was likely caused by the nonlinear feedback resistor. With proper adjustment of the circuit supply




Fig. 3. Detected oscilloscope outputs from one modulator from each 2×1 switching node from the center 4×8 section of the second-generation array at a data rate of 700 Mb/s. The 8 bit repetitive data pattern incident on the smart pixel receivers was "0 0 0 1 0 1 1 1." The optical powers were equal to 800 μ W per beam (2.8 pJ) for columns 3, 5, and 6, but column 4 required asymmetric powers of 800 μ W and 0 μ W.




Fig. 4. Eye diagram of a particular node operating at 700 Mb/s. Individual nodes operated with short term bit-error rates below 10^{-11} .

voltages and BER detector sampling point (in time), these nodes exhibited a short term bit error rate (BER) below 10^{-11} . Laser mode-hopping prevented a long term BER measurement.

The third column of the second-generation array was particularly interesting in that it contained only feedback limiting transistors. This circuit could operate at optical powers well below 1 μ W (although slowly) and could dynamically hold its state in the absence of light. We have previously described how the diode-clamped receiver can make use of this fact to operate more efficiently with optical inputs of short duration [9], [10]. Our measurements on this receiver show the same trend.

In Fig. 5, we show the supplied input optical energy for that receiver as a function of bit-rate for both nonreturn to zero (NRZ) and short pulsed inputs. The NRZ data is based on a BER for pseudorandom signals below 10^{-11} . For the pulsed data, we were unable to supply pseudorandom data, so the optical energies are based on visual inspection of the bit-pattern. We obtained clean patterns to 800 Mb/s and slightly




Fig. 5. Optical energy versus bit-rate for the nodes in the third column of the array for nonreturn to zero (NRZ) and return to zero (pulsed) data inputs. V_{dd} was lowered to 3 V for data at 25 Mb/s.

degraded patterns to 1 Gb/s. We believe the speed in both chips was limited by the driver and multiplexer circuitry.

REFERENCES

- [1] D. A. B. Miller, M. D. Feuer, T. Y. Chang, S. C. Shunk, J. E. Henry, D. J. Burrows, and D. S. Chemla, "Field-effect transistor self-electrooptic effect device: Integrated photodiode, quantum well modulator and transistor," *IEEE Photon. Technol. Lett.*, vol. 1, pp. 61–64, 1989.
- [2] L. A. D'Asaro, L. M. F. Chirovsky, E. J. Laskowski, S. S. Pei, T. K. Woodward, A. L. Lentine, R. E. Leibenguth, M. W. Focht, J. M. Freund, G. D. Guth, and L. E. Smith, "Batch fabrication and operation of GaAs-AlGaAs field effect transistor self electro-optic effect device (FET-SEED's) smart pixel arrays," *IEEE J. Quantum Electron.*, vol. 29, pp. 670–677, 1993.
- [3] J. Wieland, H. Melchoir, M. Q. Kearly, C. Morris, A. J. Mosley, M. G. Goodwin, and R. C. Goodfellow, "Optical receiver array in silicon bipolar technology with self aligned, low parasitic III–V detectors for DC-1 Gb/s parallel data links," *Electron. Lett.*, vol. 27, p. 2211, 1991.
- [4] C. Camperi-Ginestet, M. Hargis, N. Jokerst, and M. Allen, "Alignable epitaxial lift-off of GaAs materials with selective deposition using polyimide diaphragms," *IEEE Photon. Technol. Lett.*, vol. 3, p. 1123, 1991.
- [5] K. W. Goossen, J. A. Walker, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller, "GaAs MQW modulators integrated with silicon CMOS," *IEEE Photon. Technol. Lett.*, vol. 7, pp. 360–362, 1995.
- [6] T. L. Worchesky, R. J. Ritter, R. Martin, and B. Lane, "Large arrays of spatial light modulators hybridized to silicon integrated circuits," *OSA Topical Meeting on Spatial Light Modulators*, Salt Lake City, UT 1995, Postdeadline paper PD3.
- [7] A. L. Lentine, R. A. Novotny, T. J. Cloonan, L. M. F. Chirovsky, L. A. D'Asaro, G. Livescu, S. Hui, M. W. Focht, J. M. Freund, G. D. Guth, R. E. Leibenguth, K. G. Glogovsky, and T. K. Woodward, " 4×4 arrays of FET-SEED embedded control 2×1 optoelectronic switching nodes with electrical fan-out," *IEEE Photon. Technol. Lett.*, vol. 6, pp. 1126–1129, 1994.
- [8] A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, T. K. Woodward, J. E. Ford, G. F. Aplin, L. A. D'Asaro, S. P. Hui, B. Tseng, R. E. Leibenguth, D. Kossives, D. Dahringer, L. M. F. Chirovsky, and D. A. B. Miller, "3-D integration of MQW modulators over active submicron CMOS circuits: 375 Mb/s transimpedance receiver-transmitter circuit," *IEEE Photon. Technol. Lett.*, vol. 7, no. 11, pp. 1288–1290, 1995.
- [9] A. L. Lentine, L. M. F. Chirovsky, and T. K. Woodward, "Optical energy considerations for diode-clamped smart pixel optical receivers," *IEEE J. Quantum Electron.*, vol. 30, pp. 1167–1171, 1995.
- [10] T. K. Woodward, A. L. Lentine, and L. M. F. Chirovsky, "Experimental sensitivity studies of diode-clamped FET-SEED smart pixel optical receivers," *IEEE J. Quantum Electron.*, pp. 2319–2324, 1994.

Design and demonstration of a high-speed, multichannel, optical-sampling oscilloscope

Rick L. Morrison, Steven G. Johnson, Anthony L. Lentine, and Wayne H. Knox

Free-space digital optical systems have demonstrated the capability to provide thousands of optical connections between optoelectronic chips. This dense concentration of channels creates substantial challenges in monitoring individual connections for diagnostic purposes without compromising performance. From the concept of stroboscopic techniques, we have designed and constructed a multichannel optical diagnostic tool that operates analogously to an electronic-sampling oscilloscope. The tool is economically constructed by the use of commercially available video cameras and video-enhanced personal computers. An integrated software application operates the tool and displays multiple-channel waveforms. We demonstrate the oscilloscope-sampling optical waveforms of a two-dimensional optoelectronic modulator array operating at data rates from 0.5 to 4 Gbits/s. © 1996 Optical Society of America

1. Introduction

A series of system experiments has been performed to evaluate the potential and technical issues of free-space digital-optical interconnections. The basic premise is that free-space optical interconnections generated normal to the electronic component surface will provide a beneficial dense, parallel, high-speed information transfer at the chip-to-chip level. The great density of optical interconnections is achieved when the high spatial resolution of lenses that form the optical-relay framework is exploited. External electronic connectivity in these systems is generally limited to a small number of low-bandwidth signals.

The fundamental operation of one class of free-space optical systems¹⁻³ relies on the absorption of light within multiple-quantum-well (MQW) modulators in devices such as the self-electro-optic-effect device. Large arrays of laser beams are generated by diffractive elements and imaged onto arrays of modulators integrated on an optoelectronic chip. These readout beams are individually intensity

modulated during reflection to encode information processed within the isolated smart-pixel cells. The modulated beams are collected and routed to the next optoelectronic chip by a link stage that defines an interconnection fabric. Because of this method of image relay, the optical channels are isolated from each other only in the vicinity of the modulators and the detectors.

The testing and the analysis of prototype optical systems are integral to characterizing the performance and reliability of free-space photonic technology. However, this technique of transmitting information poses a serious challenge to sampling diagnostic signals within large-scale digital-photonic systems. In this paper we present a novel tool, constructed from cost-effective video and computer platforms, that utilizes the inherent optical-channel format to simultaneously monitor a two-dimensional array operating at high data rates. We demonstrate the system's effectiveness by presenting results from an electrically addressed MQW array⁴ operated at gigabits per second (Gb/s) data rates.

In Section 2 we outline the method that we implemented to monitor an array of optical channels encoded by modulator devices. Next, we describe the hardware that acts as the probe to collect and digitize the optical signals. Then we present the details of the software application developed to identify and display the waveforms. Finally, we discuss the results of a high data rate demonstration.

R. L. Morrison and A. L. Lentine are with AT&T Bell Laboratories, Naperville, Illinois 60566; S. G. Johnson is with the Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; W. H. Knox is with AT&T Bell Laboratories, Holmdel, New Jersey 07733.

Received 14 June 1995; revised manuscript received 30 October 1995.

0003-6935/96/081187-08\$06.00/0

© 1996 Optical Society of America

2. Diagnostic Method

Although information exists in both electronic and optical formats within the referenced photonic systems, the difficulty of attaining diagnostic signals becomes apparent if we examine a few schemes for data sampling. First, conventional electronic techniques are not appropriate in the typical free-space environment. Electronic contact probes must be excluded from the volume above the chip surface because they would obscure numerous optical channels. In addition, high-speed electronic test leads connected to the boundary of the electronic chip are limited in the number of sample points that can be accessed and would generate further undesired power consumption.

Although high-speed, electro-optic-sampling techniques have been developed to monitor individual electronic channels,^{5,6} the use of fiber contact probes would likewise obscure optical channels. Thus the most promising means of optical sampling is to somehow extract an image of the modulator light destined for a subsequent chip. Streak cameras are one means available to record the intensity evolution of an individual optical source by functionally scanning the film or CCD array past the focused image. However, these systems tend to be specialized (i.e., expensive) and are not easily extensible to large arrays of spots.

Up to now, the typical procedure for monitoring optical-intensity waveforms was by sampling the light reflected from the output modulators with a removable system viewport and forming a remote magnified image of the device array. A high-sensitivity photodetector was then sequentially aligned with each spot associated with a modulator to transform the signal to an electronic waveform that could be monitored by an electronic oscilloscope. This sampling procedure is too time consuming when many signals must be actively monitored. Such situations, though, are the norm when system components are being aligned and electronic parameters are being adjusted for optimal performance. This procedure is inadequate even when the mechanical alignment is computer automated if the number of channels is large.

An alternative means of sampling the two-dimensional image would be to build either a two-dimensional fiber bundle array connected to a set of receivers or a customized photodetector and receiver array whose physical layout matched that of the modulators. This solution, unfortunately, would be system specific and would still require a rigorous alignment process for coupling light into the small photosensitive areas of high-speed detectors. Indeed, the photodetector array could still be limited in the number of electrical connections that could be made off chip.

Over the course of characterizing a number of free-space optical systems, it became apparent to us that a useful diagnostic tool should satisfy two basic

criteria. Thus the objective of the project became to design and demonstrate a tool that

- simultaneously samples a two-dimensional array of high-speed optical channels in a cost-effective manner that requires a minimum of user attention either to align or to identify the regions to monitor,
- provides a user interface whose operation resembles that of a multichannel oscilloscope.

One potential technique of monitoring the device array is suggested by the process of aligning the optical channels during system construction. In the referenced systems,¹⁻³ an optical viewport option was provided so that a video camera could inspect the registration of beams to modulator windows at each optoelectronic chip. The entire electro-optic device array could be viewed and the intensity modulation directly observed during very low-speed (a few hertz) operation. Unfortunately, nominal system operating speeds are targeted toward hundreds of megabits per second per channel and even specialty CCD imagers are limited to a few kilohertz sampling rate. However, because video techniques offer the most suitable opportunity for sampling a two-dimensional array of optical channels, we have explored enhancements to this basic technique.

One method of increasing the temporal resolution is to shutter the CCD chip and thereby obtain brief time exposures. To provide sufficient diagnostic information, the exposure should be of the order of a fraction of a bit duration or less. Because this is of the order of a few hundred picoseconds at gigahertz operating speeds, mechanical shutters must be eliminated from consideration. Although high-speed electro-optical shutters are available, their exposures are usually no faster than a few nanoseconds, and light may not be sufficiently blocked during the off state. One further criterion is that the CCD sensor collect a sufficient amount of light for each system state. It is thus highly desirable to sample a periodic event repetitively to avoid photomultiplication techniques.

The solution to this challenging problem is indicated by analogous photographic procedures for studying high-speed mechanical systems. In these systems, the key is not to increase the shutter speed, but to provide a brilliant light source of exceedingly short duration to freeze the action. The sampling technique we have implemented is based on this concept of stroboscopic photography whereby the modulator is repeatedly illuminated for a brief interval by a high-intensity laser source, thus selectively capturing an image of a particular system state.

This basic sampling technique is illustrated in Fig. 1. At the top of the figure is a waveform that represents a periodic electronic signal used to modulate the absorption of a MQW device. It has a period of T_p and a bit width of T_b . Below this, an optical strobe pulse is synchronized to illuminate the




Fig. 1. General sampling scheme. The periodic electronic waveform modifies the absorption of the MQW modulator. An advancing optical-sampling pulse samples a slightly different time interval during each collection. In this example, the full waveform would be scanned in eight sampling pulses.

full device array about once per period. The reflected light is then imaged onto a low-speed CCD detector that in operation actually samples several strobe pulses. The period of the strobe is $n \cdot T_p + \Delta$, where n is an integer value, so that a waveform period is fully scanned during a time span of $n \cdot (T_p)^2 / \Delta$. The strobe offset Δ is adjusted so that the advance of the strobe is typically T_b or less during a video frame. When T_s and Δ are selected to be $\ll T_b$, the strobe pulse is able to resolve the fine scale time evolution of the waveform.

This stroboscopic method is an ideal match for modulator-based photonic systems. Within the referenced systems, spot arrays are imaged onto modulators in which electronic signals modify the optical absorption to encode information for the subsequent optoelectronic device array. The time evolution of the modulators' absorption is investigated by the replacement of the normal readout beams with a set of synchronized pulsed readout beams (serving as strobes) to scan slowly through a repeated pattern embedded in the data stream. Light is then extracted by the viewport, and the resultant remote image is viewed by a video camera. Because the acquisition rate of a standard video system is ~ 30 frames per second, the intensity evolution of an array of waveforms can be captured in a time scale of seconds.

With the diagnostic information embedded in the video signal, it next becomes necessary to digitize and process each image. If the CCD camera generates a video signal compatible with commercial standards, video digitizing boards for personal computers can be used to analyze the image. Once the video frames are digitized and stored in computer memory, the user identifies regions of interest to track intensity variation. An application monitors the optical channels and displays the intensity change over time as an oscilloscope trace. One additional feature of using a standard video to monitor the system is that the signal can be stored by video tape recorders and later reanalyzed.

3. Hardware

The project hardware was designed so that connecting it to the free-space photonic system would disturb the system as little as possible. The key functions of the hardware for the multichannel

optical oscilloscope are to

- ensure that periodic data streams modulate the optical absorption of the modulators arrays,
- synchronize short-duration stroboscopic pulses to scan the time evolution of the array slowly,
- extract and filter the reflected light and form a magnified image on a CCD camera,
- digitize the resultant video signal and store it in computer memory for analysis.

The primary components of the multichannel, optical oscilloscope are shown in Fig. 2. Typically, electronic components serve to synchronize the data stream and the stroboscopic pulse, although optomechanical methods may also be employed. Optical and video components are responsible for extracting and digitizing the image of each system state. The computer provides the platform for the software application that controls, analyzes, and displays the intensity waveforms. The nature of this application interface is discussed in Section 4.

The object to be examined is a high-speed, optoelectronic processing circuit with integrated MQW modulators. The modulators in the referenced systems have been designed to be interrogated at a wavelength of 850 nm and to operate at hundreds of megabits per second. Modulator windows in general can range in size from under 10 μm to several tens of micrometers on a side. The size is dependent on both optical and electronic performance considerations. A data generator module in the diagnostic system is either responsible for generating signals that are routed directly to the modulator or else coordinates the activity of each smart-pixel cell so that a periodic pattern persists.

The image of the modulator array is extracted by a viewport that is either added to the system as necessary or that forms part of the framework. In the referenced systems, the viewport's magnetized

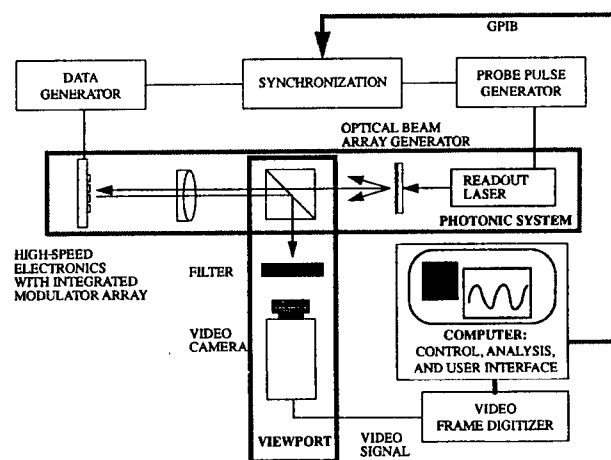


Fig. 2. Components of the optical oscilloscope: the photonic system under investigation, viewport, signal synchronization, and computer control and analysis. GPIB, general-purpose interface bus.

base allows it to be rigidly located in reserved areas of the system's steel baseplate, providing quick and easy insertion and removal. The viewport is composed of a mirror or partially reflecting beam splitter, the video camera, and potentially an illuminator. The camera's objective lens focal length determines the area of the device that will be monitored. Most monochrome video cameras are sensitive in the near infrared, although some cameras may require that a manufacturer's infrared filter placed near the focal plane be removed. To reduce the sensitivity of the camera to spontaneous laser emission between strobe pulses and light from secondary illumination, a narrow optical bandwidth filter can be employed during measurements. It is often necessary to provide additional neutral-density filters to reduce the light to a level tolerable by the CCD sensors.

Two methods can be used to produce the stroboscopic pulses. In the first method, the optical strobe pulse is generated by an external laser source linked with the viewport accessory. When the probe source is part of the viewport it is better if it provides either broad-area illumination or large spots so that only coarse registration of the illumination and modulators is required. Also, the broad-area illumination provides a simple means of identifying landmarks on the chip or of investigating local variations across large-area modulators. In this case the normal readout beams must be blocked or disabled. A light-emitting-diode source is generally not acceptable as a strobe source because of the restricted wavelength range of the quantum-well modulators and the difficulty of generating a pulse of sufficiently narrow time duration. The chief advantage of this method is that the diagnostic tool is independent of the photonic system except for the viewport tool and a clock signal shared with the system to synchronize the pulse. Also, the short-duration optical pulse may be generated by a laser more suitable for this task. In essence, this technique is suitable for all classes of electronic circuits wherein modulators are integrated solely as a means of obtaining diagnostic information.

A second method is to use the normal system readout laser to generate the probe optical pulse by supplying a new set of electronic pulse signals. Under normal system operation, this laser would generate an uninterrupted, intensity-modulated, square wave that is synchronized with each data bit as it is presented at the modulator. After modification, diagnostic readout is accomplished through disconnecting the normal clocking electronic signal and replacing it with the synchronized pulse signal. The advantages of this scheme are the ability to rely on the system's optical power sources and the ability to examine problems potentially associated with these lasers. Also, photocurrents produced in the modulators during testing may be more characteristic of those generated during normal system operation.

The low-speed responsivity of the CCD sensor has

a directly impact on the operation of the diagnostic tool. Like the standard digital-electronic-sampling oscilloscope, the modulator's absorption waveform is selectively sampled for short intervals over the course of several repeated patterns. However, because the CCD camera acts as an integrator, collecting the light throughout a video frame of 1/30-s duration, waveforms can require several seconds for accumulating. In addition, because the CCD camera is susceptible to extraneous light sources and inherent electronic noise, it is desirable to collect several pulses throughout the video frame to maximize the signal strength. Because of this collection of multiple data points, the probe pulses must be exactly registered in time to the data throughout the video frame. In addition, the sampled points thus represent an average value for each time interval and are therefore unsuitable for determining a channel bit-error rate.

The pace of the strobe delay sequence is limited by the performance of the video acquisition and analysis system. In this implementation, the sampling speed is ~ 10 samples per second. Speeds of up to 30 frames/s would be expected in future advanced processing systems. Methods of electronically controlling the data and the strobe synchronization include using a computer-controlled delay generator triggered by a system timing signal, using a multi-channel word generator that integrates pulse and data signal functions and features programmable delays, or using two signal generators that are highly synchronized but differ in the fundamental bit frequency by only a few hertz. An alternative optomechanical means of accurately delaying the optical pulse is to use a retroreflective optical relay whereby delay is introduced when a mirror is micro-positioned along the beam path.

If the system data bit rate is given by $R = 1/T_b$, the synchronization is set so that a probe pulse is generated at a rate close to R/N , where N is the number of bits associated with the repeated bit pattern. The strobe duration should be the same as or smaller than the bit duration. The shorter the pulse duration and the slower the delay scan, the greater the information about the temporal evolution of the data waveform. If a delay generator is used, it may not be possible to generate the strobe pulse faster than a rate of a few megahertz. This is adequate, provided that the delay generator accurately triggers the pulse and so long as the number of pulses within each video frame does not vary significantly. This pulse is then slowly delayed such that it samples the entire data pattern over a period of a few seconds.

The time resolution of the optical oscilloscope is determined by three factors: the duration of the optical strobe pulse, the advance of the strobe pulse during a sampling interval, and the timing jitter of the pulse and the data signals with respect to each other. In the current demonstration discussed in Section 4, a semiconductor laser was used as a strobe

source and was synchronized by two signal generators. The optical strobe width was ~ 200 ps and the pulse advanced at a rate of ~ 33 ps per video frame for 1-Gb/s data. Thus the strobe pulse duration had a dominant effect on temporal resolution. It is possible to reduce the effect of the strobe duration by performing a deconvolution of the measured pulse waveform with the data waveform.

The time resolution of the oscilloscope can be greatly improved by the use of a mode-locked laser as the strobe illumination. With a mode-locked laser it is possible to generate subpicosecond-duration optical pulses. The strobe delay can be controlled by retroreflection of the laser beam by a mirror mounted on micropositioning stages. In this manner, subpicosecond incremental delays can be added.

Finally, the electronic video signal is connected to video digitization hardware in which the image is stored in memory and becomes available to the computer processor. Rudimentary information is, of course, available to the user, who can directly view the intensity modulation on the video display. Computer-enhanced color mapping can be added to aid in distinguishing logic states. It is, however, the processing of the stored digital images that provides a more complete analysis of the waveforms.

The intensity resolution of this tool is primarily affected by the resolution of the video digitization system and the ratio of the strobe readout state energy relative to the off state. In general, the digitized video signal from current commercially available systems has an accuracy of less than 8 bits or 256 gray-scale levels. One can effectively reduce the level of video noise by integrating the spot intensity from a region of several pixels rather than by relying on a single pixel. In certain situations, this might require the image to be slightly defocused.

To monitor the spot intensities effectively, the background power of the strobe or other illumination during the off state must be considerably less than that of the readout pulse. An estimation of the required contrast ratio for the laser can be determined from

$$P_{\text{on}}/P_{\text{off}} = \text{SNR} \cdot T_{\text{off}}/T_{\text{on}},$$

where $P_{\text{on}}/P_{\text{off}}$ is the power contrast ratio of the strobe to the background, SNR is the desired signal-to-noise ratio of the waveform, and $T_{\text{on}} = T_s$ and $T_{\text{off}} \sim T_p + \Delta$ are the strobe duration and period respectively. As an example, in the demonstration presented in Section 4, $T_s = 200$ ps and $T_p + \Delta = 16$ ns. Thus for a desired signal-to-noise ratio of 1, the contrast ratio must be ~ 80 . It is this contrast-ratio requirement that determined the need to use a narrow-band optical filter to reduce the spontaneous emission background of the semiconductor laser diode satisfactorily. It must be remembered that a large background signal will further reduce the limited contrast range of the 0 and the 1 states of the MQW modulators.

One of strongest secondary advantages of select-

ing video cameras to sample the optical waveform is the unparalleled ability to record the signal with conventional video tape recorders. In this manner, system operation can be reviewed or archived to provide comparisons against future performance.

In summary, the electronic module synchronizes the data stream and the strobe pulses. The module may be as simple as a delay generator triggered by the photonic system electronics or as elaborate as two matched signal generators whose base frequencies differ by ~ 1 Hz. The strobe itself may consist of the readout lasers integrated into the system or separate lasers that form part of the viewport. The optical channels are sampled by a viewport designed to form an image on a standard video camera. The video signal is then digitized by a video-enhanced personal computer for further analysis. Aside from the custom framework needed to attach the viewport to the system, all the diagnostic hardware are readily available and are economically priced.

4. Software Application

The duty of the software application is to

- control the synchronization of the probe pulse and high-speed data stream when necessary,
- manage the sampling and analysis of the video signal,
- allow the user to identify easily the pixel regions of the video frame to monitor,
- display the time sequential intensity evolution of an array of optical channels in a manner reminiscent of a standard electronic oscilloscope.

One project objective was to select a cost-effective personal computer platform for which video acquisition hardware and image analysis software could be easily integrated in a package that could be adapted to a broad class of systems. An Apple Macintosh Quadra 840AV served as the application platform for developing and running the software for the oscilloscope interface. This system was selected on the basis of the integrated video acquisition hardware and the QuickTime video event manager software toolkit. The QuickTime system standard allows the application to be easily transferred to similar systems that adhere to this standard. As proof, the optical oscilloscope has been demonstrated on the Power Macintosh AV platform without modification. In practice, alternative platforms will also provide a suitable environment for implementing this tool.

The application software can be viewed as three basic modules: the video acquisition and analysis module, the oscilloscope control and display module, and the signal synchronization module. Each module provides a user interface for adjusting parameters and options. The software for this project was coded with the Symantec ThinkC++ compiler and relied on the Think Class libraries. Visual Architect also aided in developing a code for the user interface. Both the video and oscilloscope module

controls can be accessed through menu items permanently positioned at the top of the display screen and through popup dialogs to support individual functions interactively.

Figure 3 shows an example of the application software in operation. The leftmost window shows an image of the illuminated device array. The center video window shows the application's region selection interface. Sixteen modulator intensity waveforms obtained from the sampled video signal are shown in the rightmost window. For comparison, the intensity waveform of one modulator obtained from a high-speed photodetector is shown on the bottom left.

The synchronization of data stream and strobe pulse can be either controlled by the computer's communicating to a programmable delay generator by using the general-purpose interface bus or implemented by a tight coupling of the operation of the data and pulse generators. Under the computer-control scheme, a message would be sent to the delay generator after each frame capture, instructing the delay to be incremented by a fixed amount. Periodically, the generator would be instructed to restart the cycle. By allowing the computer to directly control the synchronization, the user can determine the degree of resolution or the speed of acquisition desired. We have not fully developed the synchronization control module in this project, as we were able to demonstrate the oscilloscope by using externally synchronized hardware.

The video module is responsible for digitizing, storing, and displaying video frames and extracting the intensity values from the designated regions of interest. The video digitization hardware and software are highly integrated with the workstation in this implementation. Apple QuickTime system software provided access to many features needed to control to these various functions. The Macintosh AV systems have video memory that is accessible by the video acquisition electronics, the central processor, and the graphics display controller, and thus

intensity values are easily attainable by the analysis routines.

The display sequence begins by opening the application tool that initializes the video acquisition hardware and displays a video window updated about once each second on the computer monitor. It may initially be necessary to provide broad-area illumination of the chip in order to identify the general location of specific modulators. Next, the illumination is reduced, and the pulsed light source is introduced (and aligned if necessary). It is not unusual to find that this light saturates the camera or video analog-to-digital converter so that a strong neutral-density filter must be inserted to reduce the intensity.

The video tool interface permits creation and manipulation of regions of interest in the video display window. The resizable video window will display either live video or a single captured image so that the user may identify the regions to monitor. A captured image is sometimes favored to avoid aligning during bits intervals when the intensity is low. The user selects the "Click Creates Region" item from the application menu to begin identifying areas to monitor. Using an interactive cursor, the user either selects an arbitrarily distributed set of regions by clicking on the center of each region or defines an array of regularly spaced regions by selecting three corners and providing the number of rows and columns. The region size can be adjusted from a single pixel to an arbitrarily sized rectangle of pixels. In the case of multiple-pixels regions, the average region intensity is calculated. The advantage of specifying a multipixel over a single-pixel region is that the alignment sensitivity is reduced and the averaging reduces some of the inherent video noise. To aid the user in accurately locating the region, a zoom feature will display a magnified region surrounding the selection point.

Another option provided is the ability to pause the video window update during analysis. By choosing to pause, the processor is able to devote a greater fraction of the time to the oscilloscope module and thereby increase its sample analysis rate. During operation, we have demonstrated the ability to analyze ~ 10 frames per second.

Once the regions are identified, the user selects the menu item "Graph Selected Regions" to create the oscilloscope traces. The oscilloscope interface is designed to present the intensity waveforms in a manner similar to that of a high-speed, multichannel oscilloscope. The waveforms can be displayed as an array of scan plots or overlaid on a common graph. Figures 3 and 4 show an example of an array of scan plots. All scans are simultaneously updated at ~ 10 points per second per channel. Once the trace has traveled across the plot, it is erased and a new trace is started. The user may choose to stop the scan at any point to examine the waveforms more closely and store the data in a file. Color is also used to




Fig. 3. Oscilloscope display interface showing an image of the illuminated modulator array, a video window for selecting the regions of interest, and the intensity traces for the designated devices.




Fig. 4. Array of scope traces taken from 16 modulators. Vertical pairs of scans show complementary data. The data scans are 1-Gb/s nonreturn-to-zero data patterns. The high-frequency scans are 1-GHz square waves.

identify waveforms in the overlay plots, and a single waveform, on being selected, can be correlated with its region in the video window. The time scale and the intensity axis of the scan region are user adjustable. Autoscaling, triggering, and data storage functions are also provided. The user is also provided with a means of defining a signal mask and selecting a specific channel for triggering the scan event.

A video signal of the operating system was collected by a video tape recorder and analyzed by the optical oscilloscope, illustrating a means of storing diagnostics for later analysis. In this mode, the synchronization module is unnecessary.

5. Demonstration

To demonstrate the capabilities of the multichannel optical sampling oscilloscope, a 2×4 array of independent electrically driven, differential modulators⁴ was monitored while operating at Gb/s data rates. Figure 5 shows an image of the modulator array in which the readout beams have been aligned to a set of the circular modulator windows. The synchronization between the data signals and the probe pulse was fixed by two high-precision, frequency-stabilized analog signal generators synchronized to a common clock to trigger digital data and pulse generators. The frequency of one generator could be adjusted to 1 part in 10^9 .




Fig. 5. Image of MQW modulator array illuminated by a beam array captured by video camera.

For the data collected in Fig. 4, the data generator was triggered at a rate of 1,000,000,001 Hz, while the probe pulse operated at frequency of 62,500,000 Hz such that the probe pulse monitored every 16th bit. Four of the differential modulators were driven by a data generator (16-bit patterns) at 1 Gb/s, and four were driven by 1-GHz square waves (i.e., a 2-Gb/s 1010 bit pattern). The voltage on the modulators was set to a 3.3-V swing that, coupled with the shift in operating wavelength caused by heating from nearby 50- Ω terminating resistors, led to a poor contrast ratio between on and off states. When the probe pulse is scanned through the data pattern at a rate of ~ 1 bit per second and sampled ~ 10 times per second, the sample spacing on the optical oscilloscope is ~ 100 ps.

Signals were collected to illustrate the similarity of the waveforms obtained by a conventional high-speed photodetector and the optical oscilloscope. Figure 6 is the intensity trace from a high-speed photodetector sampled by an electronic oscilloscope for the data pattern of 1011000 operating at a 1-Gb/s data rate. The resolution of the optical oscilloscope is limited by the optical strobe pulse. Figure 7 shows the strobe pulse intensity profile with a width of ~ 200 ps as measured by a high-speed photodetector.




Fig. 6. Data pattern of 10111000 at 1 Gb/s sampled with a conventional high-speed photodetector.

The modulators were operated at data rates from 0.5 to 4 Gb/s. The upper limit was set by the available signal generator. Throughout these data rates, the relative characteristics of the oscilloscope traces remained unchanged. The 1- and 2-Gb/s data have been presented in Fig. 4 because they show more sharply defined edges than the higher rate waveforms do. In each of the 16 traces in the figure, a common intensity and time scale was used. On close observation, it can be seen that scans can be paired vertically as a data stream and its complement. This is as expected because the electronic signal drove the center connection of the serially connected biased self-electro-optic-effect-device modulator pair. The 1-GHz square waves represent an effective 2-Gb/s bit stream that has a 1010... pattern. The patterns in the top half are 1-Gb/s 16-bit nonreturn-to-zero data streams of 1011100010111000, whereas the patterns in the bottom half are 1-Gb/s streams of 1000000111111000. The reduced size of the two signals in the right-hand column appears to be related to the modulator and not to the diagnostic tool.

The strobe spots were also severely defocused to determine whether the oscilloscope performed equivalently when a broad-area illumination was provided. Again the appearance of the traces closely




Fig. 7. Intensity evolution of optical strobe pulse as measured by a high-speed photodetector.

resembled those of Fig. 4. The system also demonstrated the ability to analyze a recorded video signal, although the trace appearance was slightly degraded. This appears to be due to the addition of video noise during the recording and a less stable horizontal and vertical frame-to-frame synchronization. Although only 16 modulators were available for this test, it was demonstrated that the tool could successfully monitor and analyze a 16×16 region without degrading performance.


6. Summary

The high concentration and overlap of information channels transmitted through a digital free-space photonic architecture preclude the use of local electronic and optoelectronic diagnostic probes. It is, however, possible to insert a viewport into the system to form a remote image of the optoelectronic device array. Through the use of stroboscopic light pulses synchronized to the data stream, it has been shown that high-speed modulator absorption can be monitored by cost-effective video cameras. The high-speed operation of free-space photonic systems is easily monitored with this novel diagnostic tool. Its chief advantage is the ability to process several optically sampled channels operating at multigigahertz rates in parallel. We have demonstrated its operation by simultaneously monitoring 16 modulators operating at data rates ranging from 0.5 to 4 Gb/s.

This work was partially sponsored by the Advanced Research Project Agency under the U.S. Air Force Rome Laboratory contract number F30602-93-C-0166.

References

1. F. B. McCormick, T. J. Cloonan, F. A. P. Tooley, A. L. Lentine, J. M. Sasian, J. L. Brubaker, R. L. Morrison, S. L. Walker, R. J. Crisci, R. A. Novotny, S. J. Hinterlong, H. S. Hinton, and E. Kerbis, "Six-stage digital free-space optical switching network using symmetric self-electro-optic-effect devices," *Appl. Opt.* **32**, 5153-5171 (1993).
2. F. B. McCormick, T. J. Cloonan, A. L. Lentine, J. M. Sasian, R. L. Morrison, M. G. Beckman, S. L. Walker, M. J. Wojcik, S. J. Hinterlong, R. J. Crisci, R. A. Novotny, and H. S. Hinton, "Five-stage free-space optical switching network with field-effect transistor self-electro-optic effect-device smart-pixel arrays," *Appl. Optics* **33**, 1601-1618 (1994).
3. F. B. McCormick, A. L. Lentine, R. L. Morrison, J. M. Sasian, T. J. Cloonan, R. A. Novotny, M. G. Beckman, M. J. Wojcik, S. J. Hinterlong, and D. B. Buchholz, "155 Mb/s operation of a FET-SEED free-space switching network," *Photon. Technol. Lett.* **6**, 1479-1481 (1994).
4. A. L. Lentine, L. M. F. Chirovsky, L. A. D'Asaro, R. F. Kopf, and J. M. Kuo, "High speed 2×4 array of differential quantum well modulators," *Photon. Technol. Lett.* **2**, 477-480 (1990).
5. J. A. Valdmanis and G. Mourou, "Subpicosecond electrooptic sampling: principles and applications," *IEEE J. Quantum Electron.* **QE-22**, 69-78 (1986).
6. B. H. Kolner and D. M. Bloom, "Electrooptic sampling in GaAs integrated circuits," *IEEE J. Quantum Electron.* **QE-22**, 79-94 (1986).



(PRIOR ART)
FIG. 1

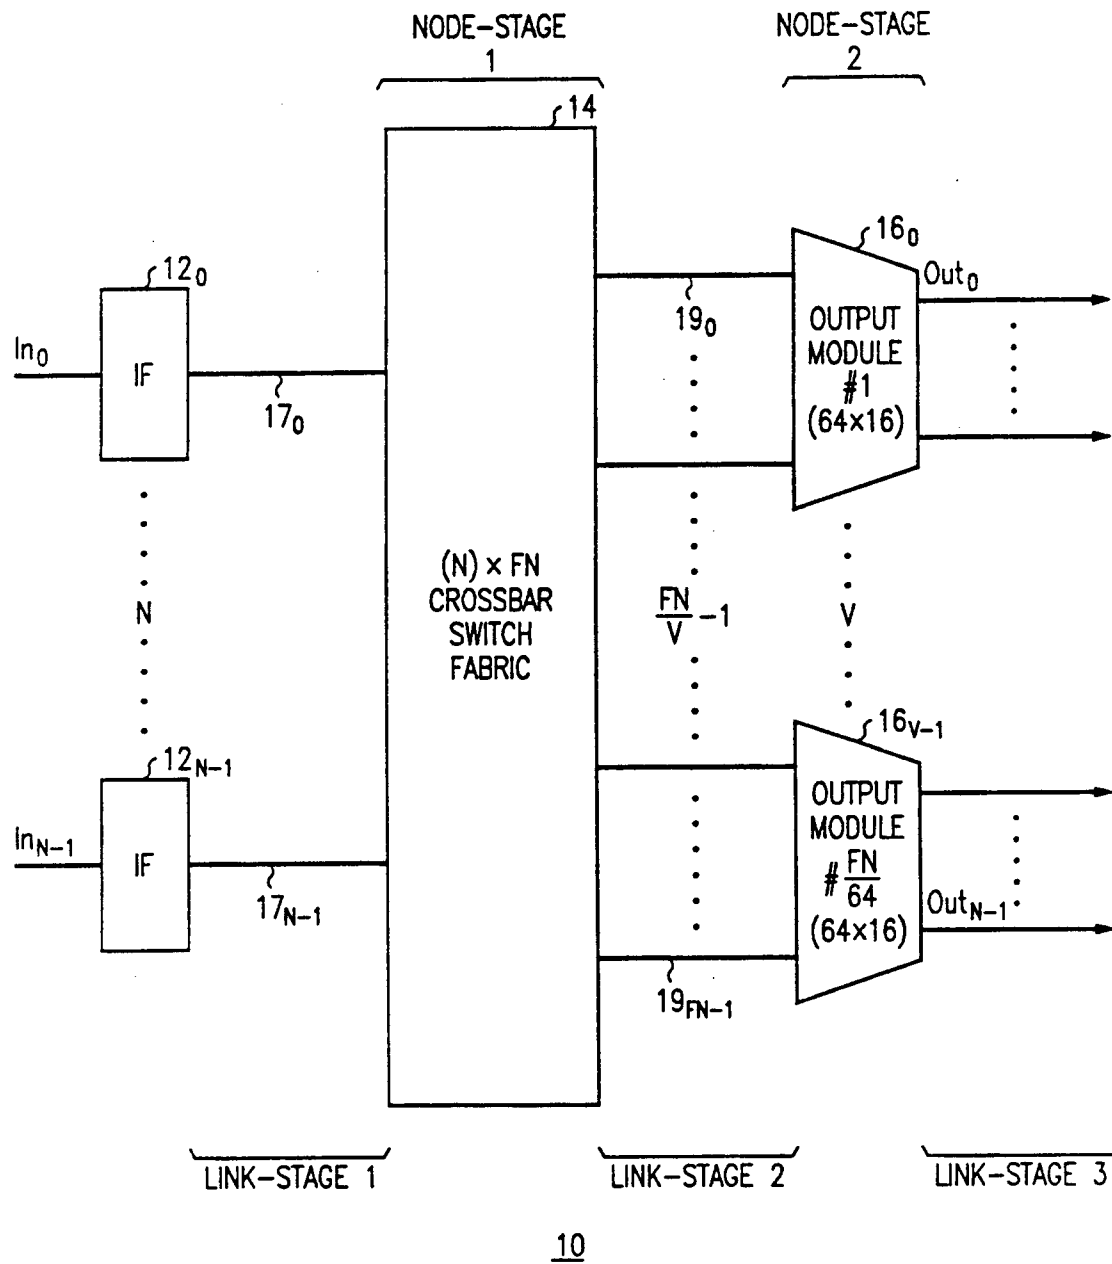


FIG. 2




FIG. 3




FIG. 4

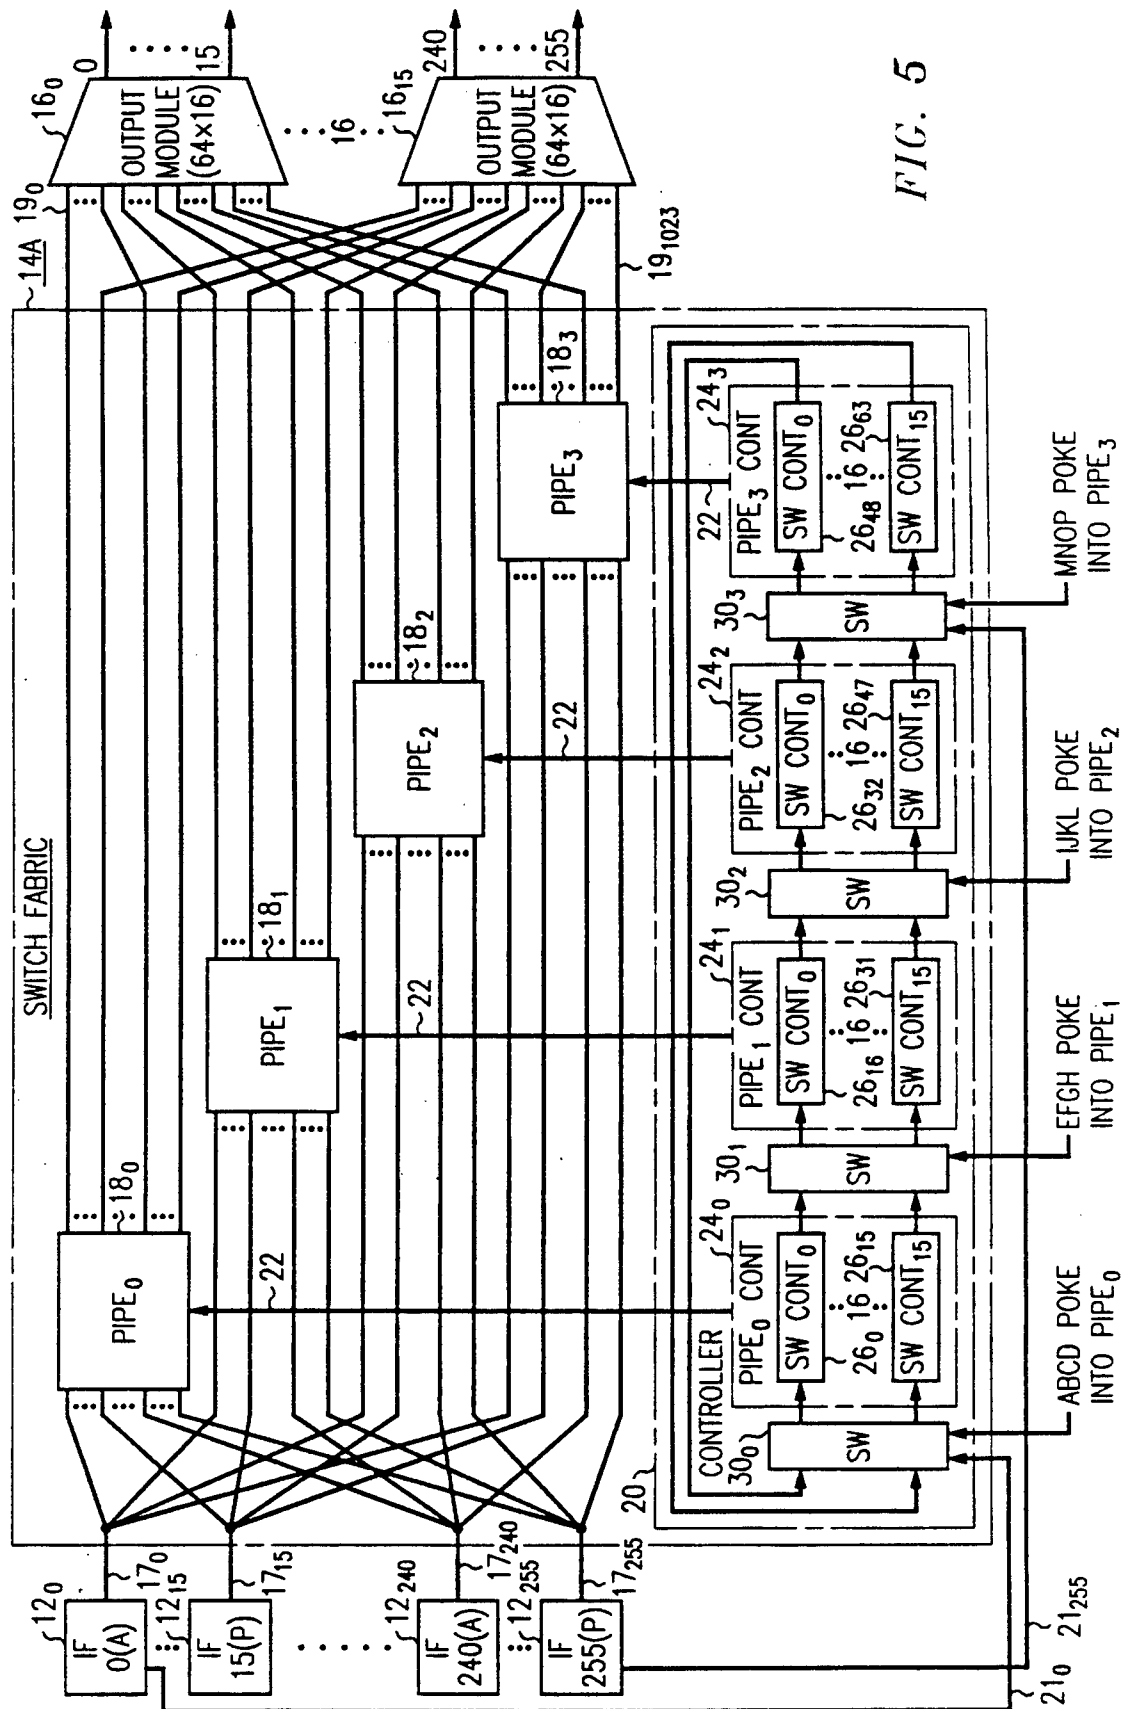


FIG. 5

TIME	PIPE ₀ CONTROLLER	PIPE ₁ CONTROLLER	PIPE ₂ CONTROLLER	PIPE ₃ CONTROLLER
i, 0	GROUP 1, PERIOD i			
i, 1		GROUP 1, PERIOD i		
i, 2		GROUP 2, PERIOD i	GROUP 1, PERIOD i	
i, 3			GROUP 2, PERIOD i	GROUP 1, PERIOD i
i, 4			GROUP 3, PERIOD i	GROUP 2, PERIOD i
i, 5	GROUP 2, PERIOD i+1			GROUP 3, PERIOD i
i, 6	GROUP 3, PERIOD i+1			GROUP 4, PERIOD i
i, 7	GROUP 4, PERIOD i+1	GROUP 3, PERIOD i+1		
i, 8		GROUP 4, PERIOD i+1		
i, 9			GROUP 4, PERIOD i+1	

FIG. 6




FIG. 7




FIG. 8




FIG. 9

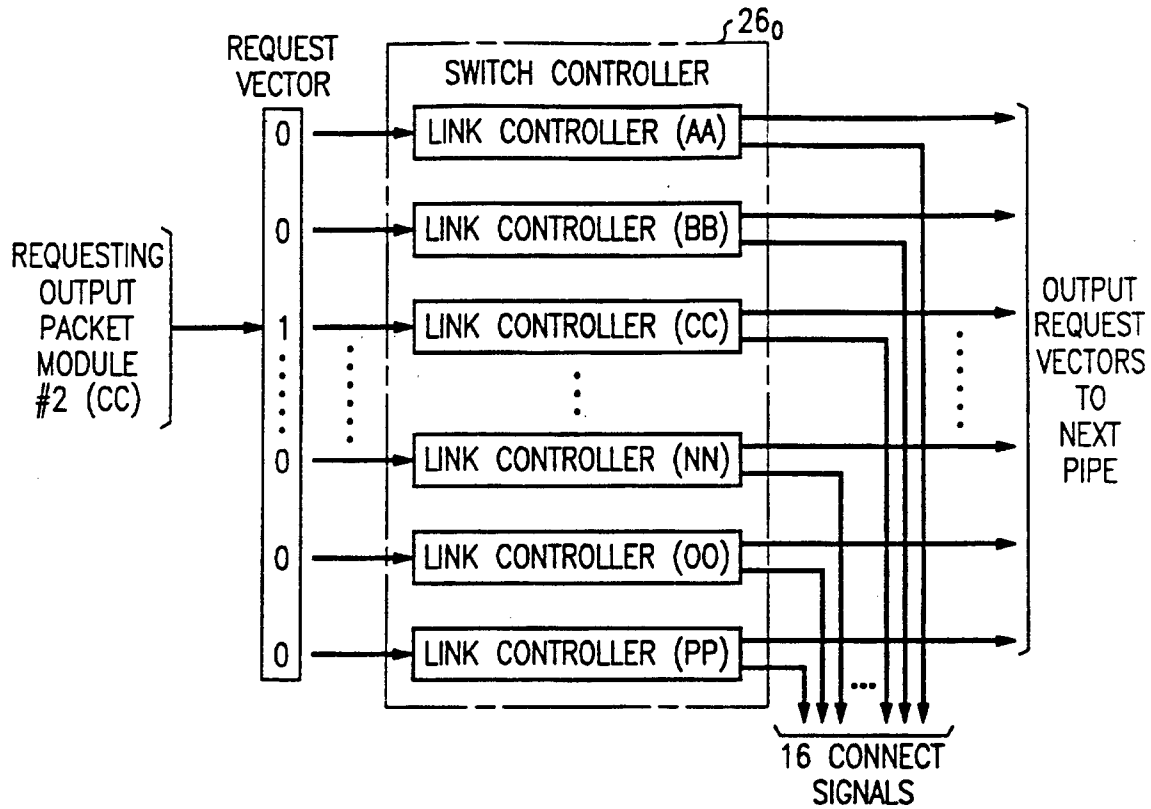


FIG. 10

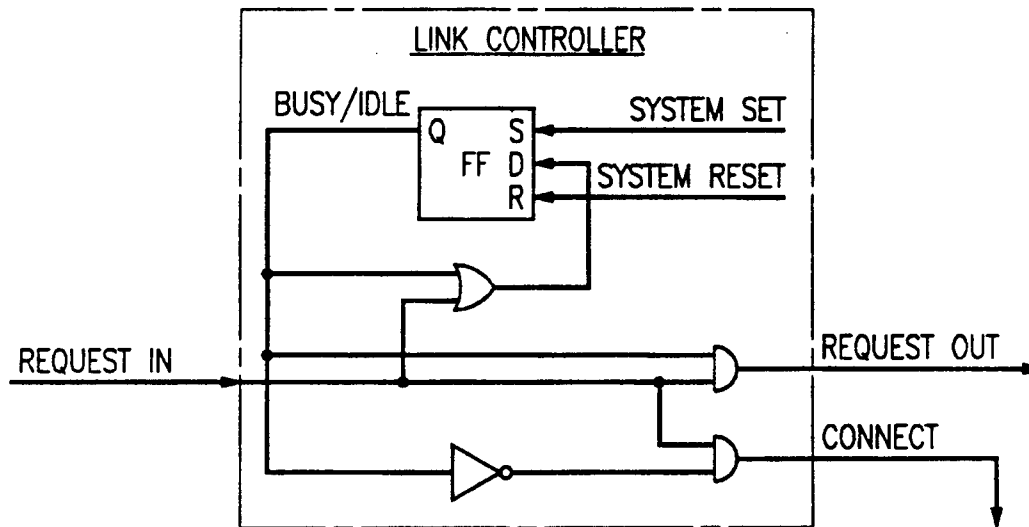


FIG. 11

FIG. 12

REQUEST IN	CURRENT BUSY/IDLE	REQUEST OUT	CONNECT	NEXT BUSY/IDLE
0	0	0	0	0
0	1	0	0	1
1	0	0	1	1
1	1	1	0	1

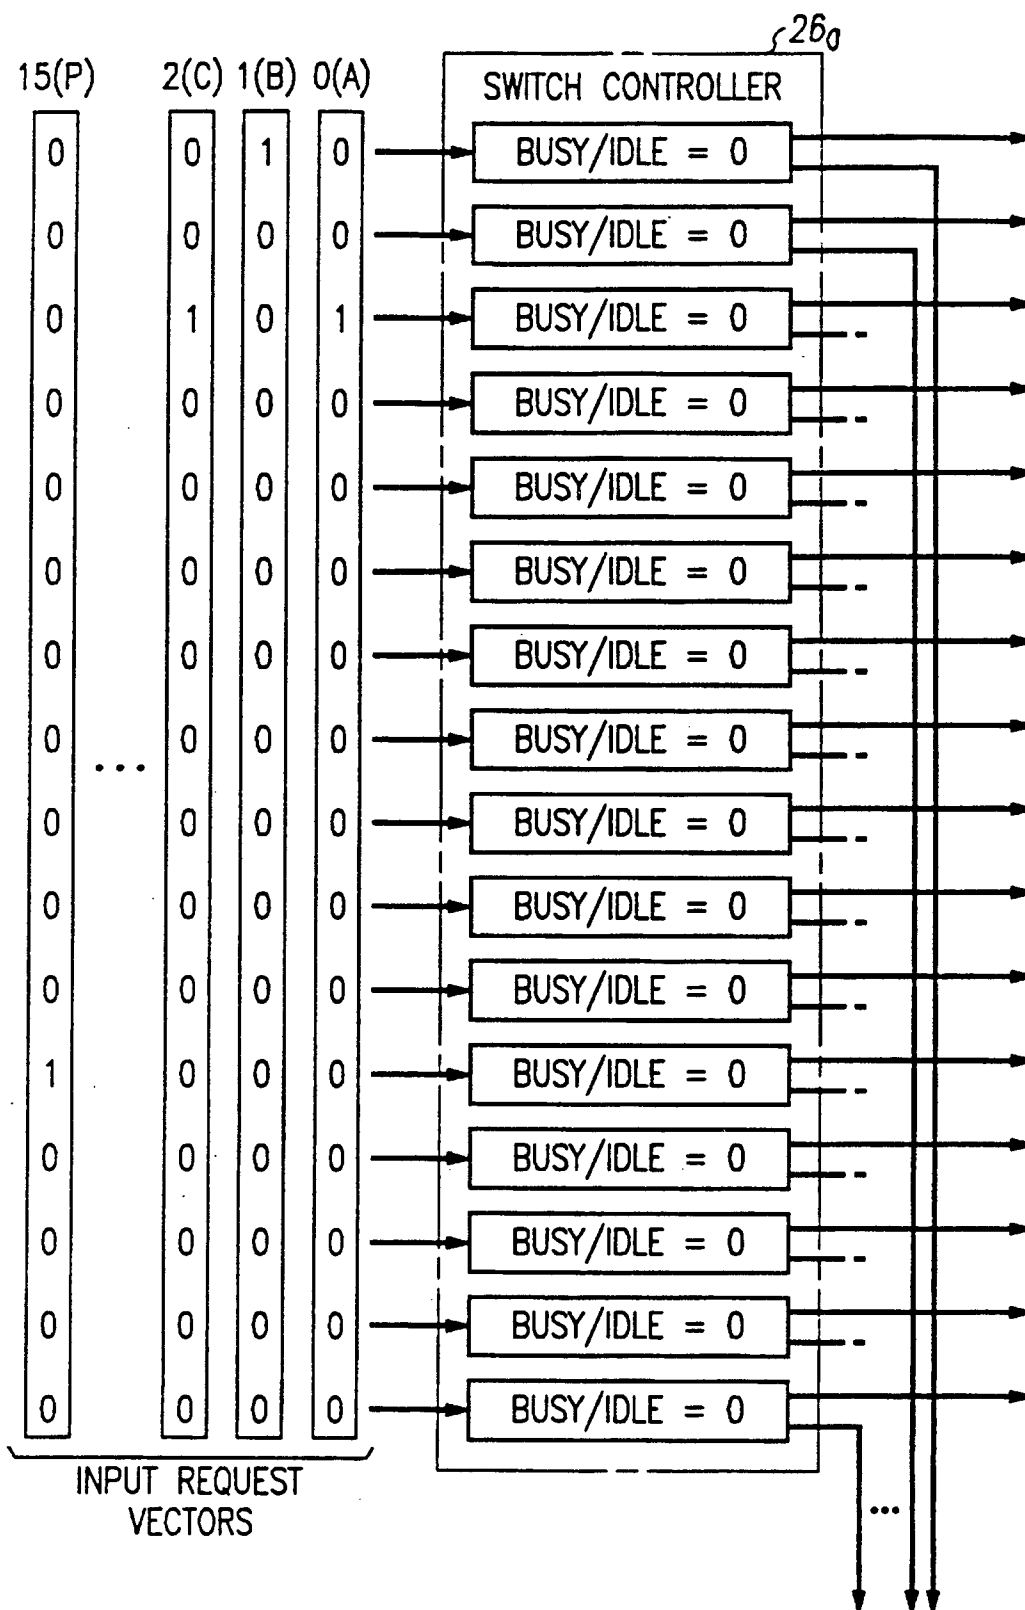


FIG. 13A




FIG. 13B




FIG. 13C





FIG. 13D

FIG. 14

TIME	CONTROLLER 24 ₀	CONTROLLER 24 ₁	CONTROLLER 24 ₂	CONTROLLER 24 ₃
i-1,17	$R_{N,i-1} \times B_{0,i}$	$R_{M,i-1} \times B_{1,i}$	* 2	$R_{0,i-1} \times B_{3,i-1}$
i-1,18	$R_{0,i-1} \times B_{0,i}$	$R_{N,i-1} \times B_{1,i}$	$R_{M,i-1} \times B_{2,i}$	$R_{P,i-1} \times B_{3,i-1}$
i-1,19	$R_{P,i-1} \times B_{0,i}$	$R_{0,i-1} \times B_{1,i}$	$R_{N,i-1} \times B_{2,i}$	
i,0	$R_{A,i} \times B_{0,i}$	$R_{P,i-1} \times B_{1,i}$	$R_{0,i-1} \times B_{2,i}$	
i,1	$R_{B,i} \times B_{0,i}$	$R_{A,i} \times B_{1,i}$	$R_{P,i-1} \times B_{2,i}$	
i,2	$R_{C,i} \times B_{0,i}$	$R_{B,i} \times B_{1,i}$	$R_{A,i} \times B_{2,i}$	* 3
i,3	$R_{D,i} \times B_{0,i}$	$R_{C,i} \times B_{1,i}$	$R_{B,i} \times B_{2,i}$	$R_{A,i} \times B_{3,i}$
i,4		$R_{D,i} \times B_{1,i}$	$R_{C,i} \times B_{2,i}$	$R_{B,i} \times B_{3,i}$
i,5		$R_{E,i} \times B_{1,i}$	$R_{D,i} \times B_{2,i}$	$R_{C,i} \times B_{3,i}$
i,6		$R_{F,i} \times B_{1,i}$	$R_{E,i} \times B_{2,i}$	$R_{D,i} \times B_{3,i}$
i,7	* 0	$R_{G,i} \times B_{1,i}$	$R_{F,i} \times B_{2,i}$	$R_{E,i} \times B_{3,i}$
i,8	$R_{E,i} \times B_{0,i+1}$	$R_{H,i} \times B_{1,i}$	$R_{G,i} \times B_{2,i}$	$R_{F,i} \times B_{3,i}$
i,9	$R_{F,i} \times B_{0,i+1}$		$R_{H,i} \times B_{2,i}$	$R_{G,i} \times B_{3,i}$
i,10	$R_{G,i} \times B_{0,i+1}$		$R_{I,i} \times B_{2,i}$	$R_{H,i} \times B_{3,i}$
i,11	$R_{H,i} \times B_{0,i+1}$		$R_{J,i} \times B_{2,i}$	$R_{I,i} \times B_{3,i}$
i,12	$R_{I,i} \times B_{0,i+1}$	* 1	$R_{K,i} \times B_{2,i}$	$R_{J,i} \times B_{3,i}$
i,13	$R_{J,i} \times B_{0,i+1}$	$R_{I,i} \times B_{1,i+1}$	$R_{L,i} \times B_{2,i}$	$R_{K,i} \times B_{3,i}$
i,14	$R_{K,i} \times B_{0,i+1}$	$R_{J,i} \times B_{1,i+1}$		$R_{L,i} \times B_{3,i}$
i,15	$R_{L,i} \times B_{0,i+1}$	$R_{K,i} \times B_{1,i+1}$		$R_{M,i} \times B_{3,i}$
i,16	$R_{M,i} \times B_{0,i+1}$	$R_{L,i} \times B_{1,i+1}$		$R_{N,i} \times B_{3,i}$
i,17	$R_{N,i} \times B_{0,i+1}$	$R_{M,i} \times B_{1,i+1}$	* 2	$R_{0,i} \times B_{3,i}$
i,18	$R_{0,i} \times B_{0,i+1}$	$R_{N,i} \times B_{1,i+1}$	$R_{M,i} \times B_{2,i+1}$	$R_{P,i} \times B_{3,i}$
i,19	$R_{P,i} \times B_{0,i+1}$	$R_{0,i} \times B_{1,i+1}$	$R_{N,i} \times B_{2,i+1}$	
i+1,0	$R_{A,i+1} \times B_{0,i+1}$	$R_{P,i} \times B_{1,i+1}$	$R_{0,i} \times B_{2,i+1}$	
i+1,1	$R_{B,i+1} \times B_{0,i+1}$	$R_{A,i+1} \times B_{1,i+1}$	$R_{P,i} \times B_{2,i+1}$	
i+1,2	$R_{C,i+1} \times B_{0,i+1}$	$R_{B,i+1} \times B_{1,i+1}$	$R_{A,i+1} \times B_{2,i+1}$	* 3
i+1,3	$R_{D,i+1} \times B_{0,i+1}$	$R_{C,i+1} \times B_{1,i+1}$	$R_{B,i+1} \times B_{2,i+1}$	$R_{A,i+1} \times B_{3,i+1}$
i+1,4		$R_{D,i+1} \times B_{1,i+1}$	$R_{C,i+1} \times B_{2,i+1}$	$R_{B,i+1} \times B_{3,i+1}$

*FIG. 15*

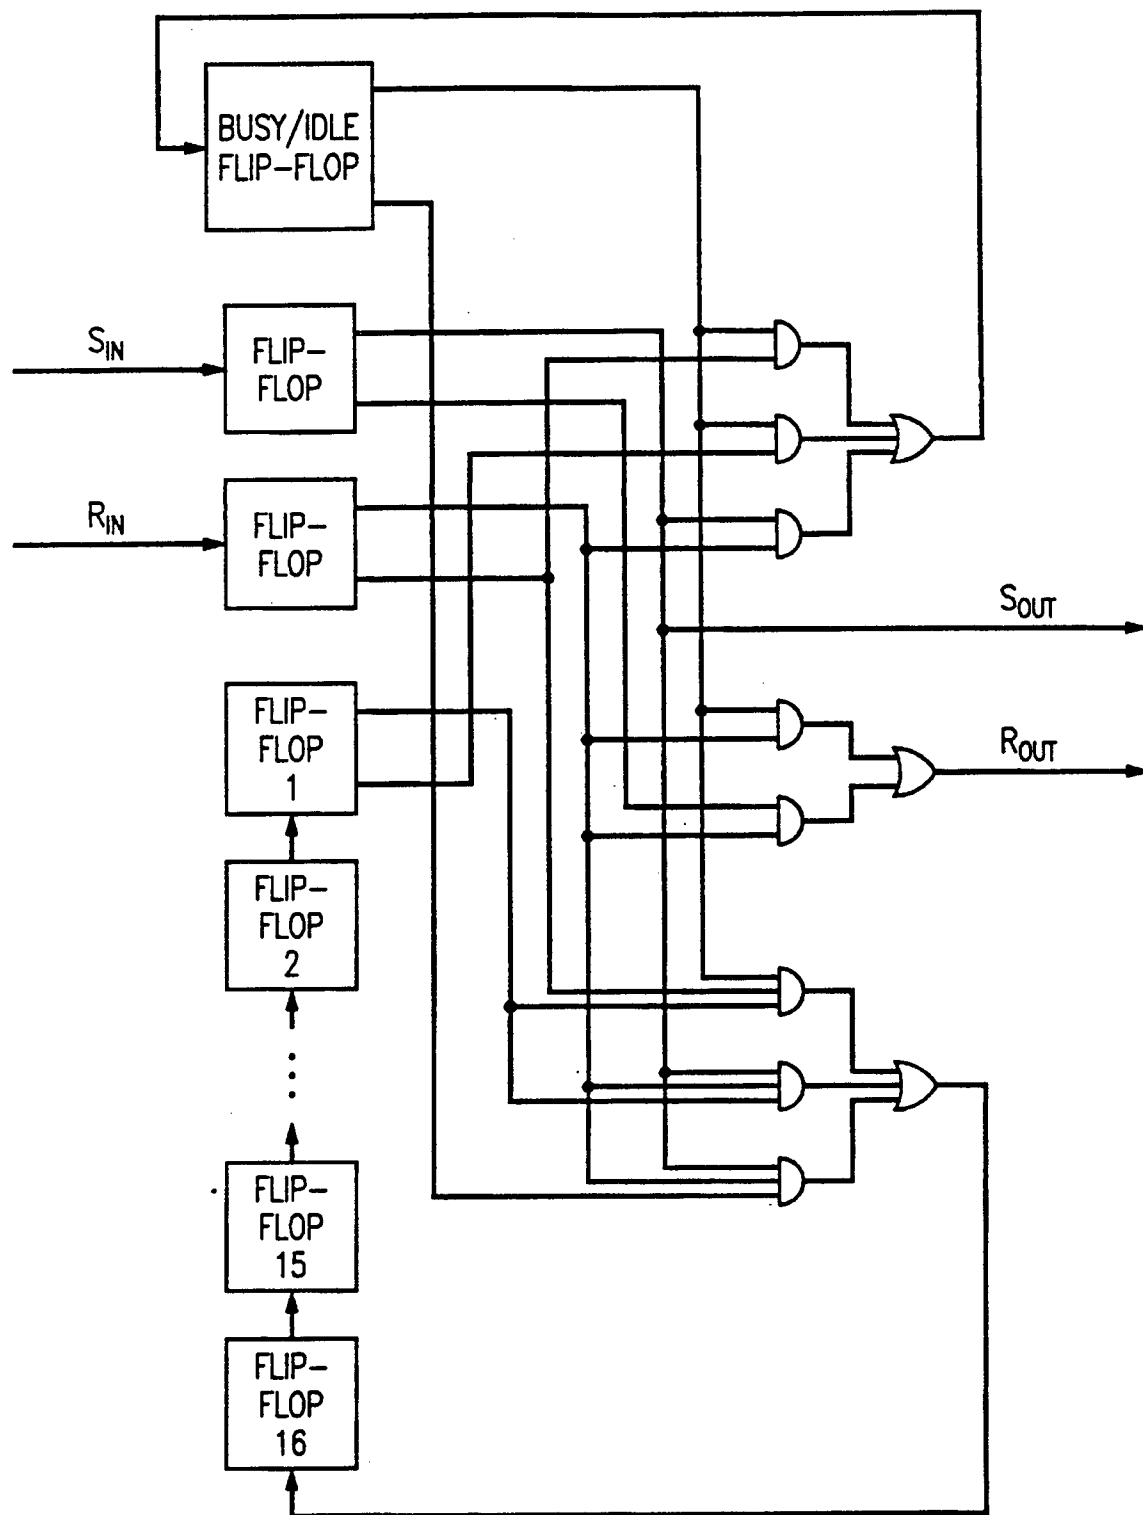


FIG. 16




FIG. 17

TERABIT PER SECOND PACKET SWITCH HAVING DISTRIBUTED OUT-OF-BAND CONTROL OF CIRCUIT AND PACKET SWITCHING COMMUNICATIONS

CROSS REFERENCES

This application is related to the following co-pending applications: "TERABIT PER SECOND PACKET SWITCH" by Thomas Cloonan and Gaylord Richards, Ser. No. 08/366,708; "TERABIT PER SECOND ATM PACKET SWITCH HAVING DISTRIBUTED OUT-OF-BAND CONTROL" by Thomas Cloonan and Gaylord Richards, Ser. No. 08/367,489; "METHOD AND APPARATUS FOR DETECTING AND PREVENTING THE COMMUNICATION OF BIT ERRORS ON A HIGH PERFORMANCE SERIAL DATA LINK" by Thomas Cloonan, Ser. No. 08/366,706; and "TERABIT PER SECOND DISTRIBUTION NETWORK" Thomas Cloonan and Gaylord Richards, Ser. No. 08/366,707; and "APPARATUS AND METHOD FOR REDUCING DATA LOSSES IN A GROWABLE PACKET SWITCH" by Thomas Cloonan and Gaylord Richards Ser. No. 08/366,705.

TECHNICAL FIELD

The invention relates to large telecommunication switches and more particularly to large telecommunication switches that use data packets in order to communicate at aggregate throughputs at the one terabit per second level.

DESCRIPTION OF THE PRIOR ART

Telecommunications have long used digital switching to encode, multiplex, transmit and decode audio frequencies in order to carry the millions of telephone voice calls of the world. Telecommunication switches for voice calls have grown to very large sizes to keep pace with the demand. Most of the switching systems that route and control voice call traffic are called circuit switches, which means that for each call a type of bi-directional audio circuit is set up between the calling party and the called party. The circuit that is set up has the bandwidth and transport timing necessary to simulate a face-to-face conversation without objectionable distortion or time delays.

An alternative to circuit switching is called packet switching. For packet switching, the calling party is responsible for converting the information into one or more packets. This information could be encoded voice, it could be encoded computer data, or it could be encoded video. The number of the called party is typically included in a packet header to guide the packet to its destination. The packet switching network has the task of routing each packet to its respective destination without undue delay. The called party usually has the equipment to receive the packets and decode the information back into an appropriate form.

The extremely rapid growth of packet switching traffic carrying voice, computer (LAN/WAN), facsimile, image and video data to an ever widening variety of locations, along with the proposals for a National Information Infrastructure, has challenged both the packet switch protocols and system architecture's.

Many vendors and service providers have joined forces to define a global standard that would permit packet switching services to be provided in a ubiquitous fashion. The result of this coordinated effort has been the rapid development and deployment of an Asynchronous Transfer Mode (ATM) as a

means of efficiently routing and transporting data packets that have stochastically-distributed arrival rates according to the recent ATM standard. ATM is thus a packet-oriented standard, but unlike most of its data packet predecessors (X.25, frame relay, etc.), ATM uses short, fixed-length, 53-byte packets that are called cells. ATM also uses a very streamlined form of error recovery and flow control relative to its predecessors. In fact, the ATM standard essentially eliminates most error protection and flow control at the link level, leaving these functions to higher level protocols at the edges of the network. This approach permits rapid routing of the short cells with minimal network delay and jitter, making ATM compatible with voice, data and video services. ATM has been embraced by the computer, LAN, and WAN industries, so a seam-less packet communication from the source computer through LANs, WANs, and the public-switched network is a reality.

If this level of connectivity becomes available to the average consumer and if advanced broadband services that combine voice, broadband data and video are similarly available at a reasonable price, then the volume of ATM traffic that may be generated in the future is virtually limitless. As a result, the number and size of the switches and cross-connects required to route this ATM packet traffic may also grow by phenomenal rates within the next decade. ATM switches and cross-connects for toll and gateway applications may require aggregate bandwidths ranging from 155 gigabits per second (1000 inputs at SONET OC-3 155 megabits per second rates) to 2.4 terabits per second (1000 inputs at SONET OC-48 2.4 gigabits per second rates). Additionally, if demand for broadband services to the home and/or LAN/WAN connectivity through the public-switched network grows as some experts believe, then local telephone exchange carriers may require ATM switches and cross-connects for metropolitan area network (MAN) applications having aggregate bandwidths ranging from 100 gigabits per second (50,000 inputs at Ethernet 10 megabits per second rates and 20 percent occupancy) to 775 gigabits per second (50,000 inputs at SONET OC-3 155 megabits per second rates and 10 percent occupancy).

By necessity, most of the current architectural research and hardware/software development for ATM switches has concentrated on switches with much smaller aggregate bandwidths that will meet the more near-term needs of the marketplace. For example, most of the proposals within the LAN/WAN community have supported aggregate bandwidths ranging from 150 megabits per second to 12 gigabits per second, and most of the published proposals within the telecommunications community have supported aggregate bandwidths ranging from 20 gigabits per second to 160 gigabits per second. Extensions of most of these architecture's to larger sizes usually produce systems that are cost prohibitive, size prohibitive, and/or physically unrealizable because of limits of the underlying technology.

For example, very common designs for large, high-throughput switches use a multi-stage interconnection network containing multiple stages of switching nodes (node-stages) interconnected by stages of links (link-stages) to provide multiple paths between input ports and output ports. Clos, Banyan and Benes networks are examples of such networks. A multiple stage network design can yield networks with very high levels of performance (low blocking probabilities, low delay, high degrees of fault tolerance, etc.), and may result in low system-level costs, because network resources (nodes and links) are time-shared by the many different paths that can be set up within the network. Physically realizing a multistage network for ATM is, however, a problem.

The design of any large, high-throughput ATM switching architecture must address two fundamental issues that profoundly effect the overall performance of the resulting ATM switch. The first of these issues is cell loss due to blocking within the internal links of the distribution network (also known as the switching fabric), and the second is cell loss due to contention for output ports by two or more ATM cells that pass through the switch at the same moment in time. The first issue can usually be solved by designing a network with sufficient switching fabric (nodes and links) so that multiple paths exist between input ports and output ports. As a result, if two or more ATM cells attempt to use the same shared resource (nodes or links) within the switching fabric, the cells can usually find two disjoint paths that eliminate the internal network blocking problem. The second issue requires the switch designer to identify a technique for handling simultaneous cells.

A general design technique for a switch to handle cells destined for the same output port is analyzed in an article, A Growable Packet Switch Architecture, IEEE Transactions on Communications, February, 1992, by Eng et al. and in another article The Knockout Switch, ISS AT&T Technical Papers, 1987, by Yeh et al. This general design technique segments the switch into two distinct parts, as shown in FIG. 1. A $N \times (FN)$ distribution network (which provides for N input ports) and a bank of K $m \times n$ output packet modules (which provide for a total of $M = Kn$ output ports). Given that each of the links emanating from the distribution network is required to be terminated at one of the inputs to an output packet module, it can be seen that the equation $FN = Km$ must be satisfied. The switching fabric is a memory-less $N \times (FN)$ fanout switch whose function is to route an arriving ATM cell to any of the m inputs on the output packet module connected to the cell's desired output port. The output packet module is a $m \times n$ switch with buffers that are available for storing cells that must be delayed when two or more cells contend for a particular output port. If the arriving traffic is uniformly distributed across all output ports and if the buffers within the output packet modules are sufficiently large, then the ratio $m:n$ can always be chosen large enough to force the cell loss probability within the network to be below any desired cell loss probability level. In fact, if the network size (N) is large and if R represents the switch loading, then the cell loss probability of a network with $m \times n$ output packet modules as shown by Eng et al is given by:

$$P(\text{cell loss}) = \left[1 - \frac{m}{(nR)} \right] \left[1 - \sum_{k=0}^m \frac{(nR)^k e^{-(nR)}}{k!} \right] + \frac{(nR)^m e^{-(nR)}}{m!}$$

Present packet switches have acceptable cell loss probabilities of approximately 10^{-12} , so any loss probability smaller than present units are considered acceptable.

Besides the ATM cell losses because of internal contentions, in an ATM Packet Switch where all of the N cells arrive simultaneously at the inputs of the distribution network, the cells must be processed by each stage of the path hunt processing pipeline before the next group of N cells arrives at the network input ports. If, for example, the incoming transmission lines support SONET OC-48 2.5 gigabits per second bit-rates, then the group of N ATM cells that arrive together must be processed and sent on to the next stage of the pipeline every 176 nano second (the duration of an ATM cell on a 2.5 gigabits per second). For large values of N , a substantial amount of processing power will therefore be required to complete the path hunt operations for all N cells. (Note: If $N=256$, then 1.45×10^9 path hunts must be completed every second, which corresponds to an average

processing rate of one path hunt every 684 pico seconds). Present commercial microprocessors can process approximately 100 million instructions per second. If each path hunt took only one instruction, at least 15 of microprocessors would spend 100% of their processing time performing these path hunts. Thus, a controller based on something other than a single microprocessor will be necessary for a large ATM packet switch.

Two approaches to solving the path hunt problem can be envisioned. One approach uses in-band, i.e. self-routing, control techniques to perform the path hunts. For in-band control techniques, the connection requests are prepended to the ATM cells and routed through the switch along the same paths used by the following ATM payload. This approach typically requires parallel processing elements to be distributed throughout all of the nodes in the network, resulting in relatively complicated hardware within each node of the network in order to perform localized path hunt calculations (on only the cells that pass through that node) when determining how to route the arriving connection requests and ATM cells. The second approach uses out-of-band control techniques whereby the controller and switch fabric are logically separated, so connection requests must be routed to the controller before the control signals resulting from the path hunt are injected into the switch fabric to set the paths. This second approach requires that the out-of-band controller have tremendous processing power, (as mentioned above), because of the many path hunt operations that must be performed in a very short period of time.

Since the path hunting operations in switches using in-band control techniques are only based on localized traffic information and not on global information regarding all of the switch traffic, the connections may not always be routed in optimal fashion. As a result, networks based on in-band control techniques often require more switch fabric (stages and nodes) to provide the same operating characteristics as a less expensive switch based on out-of-band control techniques. In addition, out-of-band control ATM switch architecture's share many similarities with the partitioning of many existing telecommunication switching and cross-connect products that have centralized controllers, so the development of a system based on such an architecture should yield fewer design problems than an architecture based on an entirely new architectural approach. Thus, ATM switch designers might consider an out-of-band control ATM switch to benefit from the lower overall hardware costs and more standard architectural partitioning. On the other hand, the difficulties associated with performing path hunts in an out-of-band controller for N arriving ATM cells is the time required by a standard partitioned out-of-band controller to perform N path hunts might influence ATM switch designers to consider an in-band control switch design. Assuming a single path hunt requires at least one read from a busy-idle memory and one write to a busy-idle memory, path hunt requires $2N$ accesses to memory. If $N=256$, then 512 memory accesses are required every 176 nano second, so the average memory access time must be 340 pico second. Since 340 pico second memories are not commercially available, a path hunt scheme different than the present standard architectural partitioning is required for any out-of band controller.

The high probability that large ATM switches will be required coupled with the uncertainties and shortcomings of present architecture's demonstrate a strong need in the art for a flexible packet switch architecture that will operate with throughputs at the terabit per second levels.

It is an object of the present invention to provide an ATM packet switch architecture that has a large aggregate bandwidth.

It is yet another object of the invention to provide an ATM packet switch that is flexible to meet current and future telecommunication needs as they evolve.

It is another object of the invention to provide an ATM packet switch that has a high degree of fault-tolerance.

SUMMARY OF THE INVENTION

Briefly stated, in accordance with one aspect of the invention, the foregoing objects are achieved by providing a switch architecture that has a single stage switch fabric which will operate equally well with circuit switched communications and with packet switched communications, or with a combination of both. Such a switch has the advantage of meeting present demands for circuit switched and packet switched communications and the flexibility to evolve as the demands for these services change, without major changes to the switch. Separate circuit switched and packet switched out-of-band controllers are required, as are separate input interfaces and output modules. But these are modular and can be easily modified as needs evolve. The central switch fabric does not need to change.

BRIEF DESCRIPTION OF THE DRAWING

FIG. 1 is a block diagram of a generalized growable packet switch.

FIG. 2 is a slightly re-drawn FIG. 1.

FIG. 3 is a block diagram of a growable packet switch in which the switch fabric is partitioned into L multiple pipes according to the present invention.

FIG. 4 is a block diagram, similar to FIG. 3, of a specific embodiment of the present invention having four pipes ($L=4$) and showing a configuration for the pipes.

FIG. 5 is a simplified block diagram of the embodiment shown in FIG. 4 which shows greater details of the controller.

FIG. 6 illustrates the timing sequences of requests to the controller shown in FIG. 5.

FIG. 7 is a simplified block diagram of an embodiment of an output module.

FIG. 8 is an illustrative example of rolling and its operation in a plan view of an amusement park and its satellite parking lots.

FIG. 9 shows plots of calculated values of various ATM cell loss probabilities both with and without the assignment of preferences.

FIG. 10 is a simplified block diagram of a representative switch controller and its link controllers.

FIG. 11 is a detailed logic diagram of a link controller.

FIG. 12 is a state table for the link controller shown in FIG. 11.

FIGS. 13A-13D when joined together show the operation of a switch controller in response to a sequence of requests.

FIG. 14 illustrates the rolling of path hunting requests through a switch having four pipe controllers according to the present invention.

FIG. 15 shows a plot of packet loss probability versus the percentage of faulty links in the switch fabric.

FIG. 16 is a detailed logic diagram of a link controller for indefinite length packets.

FIG. 17 is a block diagram of a switch that routes both circuit switched communications and packet switched communications through the same switch fabric.

DETAILED DESCRIPTION

Referring now to FIG. 2, a large, generalized switch 10 for ATM communications, is shown in block diagram form. ATM switch 10 has a number of input interfaces 12_0-12_{N-1} , a switch fabric 14, and buffered output modules 16_0-16_{N-1} . For ATM operation, input interfaces 12_0-12_{N-1} are high speed digital amplifiers that serve as a matching networks and power amplifiers for fanning out information received on their inputs to multiple input ports of the switch fabric 14. Each of the input interfaces 12_0-12_{N-1} also needs a capability to store one ATM cell, as will be explained below. Similarly for ATM operation, buffered output modules 16_0-16_{N-1} are concentrators that are buffered to reduce packet loss when two or more packets are directed to and contend for the same output of outputs Out_0-Out_{N-1} .

Switch fabric 14 includes a fanout F where each of the outputs from the input interfaces 12_0-12_{N-1} is fanned out to F inputs within switch fabric 14, such that if ATM switch 10 is an $N \times N$ switch then switch fabric 14 will have FN internal inputs and FN outputs to output modules 16_0-16_{N-1} . Output Modules 16_0-16_{N-1} have a fanin or concentration factor of F in order to convert the FN outputs of the switch fabric 14 to N output module outputs Out_0-Out_{N-1} . Each output module 16_0-16_{N-1} stores arriving ATM packets in FIFO queues, and then routes the ATM packets at the front of each of these FIFO queues to their desired outputs Out_0-Out_{N-1} when the output ports are available.

Switch fabric 14 is a general distribution network which may be a network of switches, specifically crossbar switches, to provide multiple paths from each of its input ports 17_0-17_{N-1} to each of its output ports 19_0-19_{FN-1} . However, it becomes highly impractical to make an $N \times N$ switch out of a single crossbar to operate as the switching component of switch fabric 14 when the size of N exceeds 32. Thus, some other way is needed to realize the general architecture shown in FIG. 2.

Referring now to FIG. 3, an ATM switch 10A that is both practical and possible for N inputs where the size of N is at least 256, is shown. Multiple paths from each input 17_0-17_{N-1} through the switch fabric 14A are provided to prevent blocking. These multiple paths are partitioned into groups called pipes with each pipe providing exactly one path between each input port 17_0-17_{N-1} and each output port 19_0-19_{FN-1} within the network. Thus, switch fabric 14A is made up of multiple pipes 18_0-18_{L-1} . The output modules 16_0-16_{N-1} are essentially the same as the output modules shown in FIG. 2.

Switch fabric as seen in co-pending and commonly assigned application entitled TERABIT PER SECOND DISTRIBUTION NETWORK, which is hereby incorporated by reference, is a single stage, memoryless, and non-self routing network. Since the switch fabric 14A is not unconditionally non-blocking as a full $N \times N$ crossbar switch would be, a controller 20 is included to hunt for a path through the four pipes for each ATM cell. Since each of the pipes 18_0-18_3 contains a path that could transport the ATM Cell, the real purpose of the controller 20 is to find a path that is not blocked.

For ATM switch 10A, if the number of input lines, N is equal to 256 and if each input line is operated at a standard 2.5 Gigabits per second data rate, its aggregate throughput will be 0.640 terabits per second. Scaling or growing such an ATM switch by a factor of two to 512 input lines and output lines should be straightforward and result in aggregate throughputs of greater than 1 Terabits per second. Scaling to an ATM switch size of 1024×1024 is considered within the

present technology, and the architecture of the present invention is believed to be extensible even further as the speed of commercially available components increases and as new, faster technologies are developed.

Referring now to FIG. 4, a specific embodiment of an ATM switch 10A is shown. In this specific embodiment ATM switch 10A has two hundred fifty six input interfaces 12₀-12₂₅₅ which are connected to two hundred fifty-six ATM input lines In₀-In₂₅₅. The outputs of the input interfaces are connected to the input ports 17₀-17_{N-1} of the switch fabric 14A. The switch fabric 14A contains a total of sixty-four 16×16 crossbar switches 15₀-15₆₃ which are partitioned into four pipes 18₀-18₃. The fanout F is equal to four which if the number of output ports=FN results in 1024 output ports 19₀-19₁₀₂₃. The output ports 19₀-19₁₀₂₃ are respectively connected to the inputs of sixteen 64×16 output packet modules 16₀-16₁₅. The sixteen 64×16 output packet modules are connected to two hundred fifty six outputs Out₀-Out₂₅₅. Those skilled in the art will recognize that other combinations of components could have been used. For example thirty two 32×8 output modules could have been used instead of the 64×16 output modules shown in FIG. 4.

ATM switch 10A also has a controller 20 which has the tasks of hunting and finding an available pipe through the switch fabric 14A for each ATM packet. The controller 20 uses the fact that the switch fabric 14A is partitioned into four pipes to break the pipe hunting tasks into four parallel pipe hunting tasks that are each temporally shifted by an acceptable amount. Details of one embodiment of such a controller 20 are shown in FIG. 5.

For the 0.640 Terabits per second, N=256 embodiment mentioned previously and shown in FIGS. 4 and 5, the controller 20 may be contained on approximately eight printed circuit boards. Controller 20 would accept up to 256 sixteen-bit request vectors from up to 256 line input interfaces 12₀-12₂₅₅ and perform path hunts on each of these request vectors within each 176 nanosecond. ATM cell interval to create the 1024 sixteen-bit connect vectors used to established connections within the switch fabric 14A. This requires that controller 20 operate with a processor clock rate of at least 46 megabits per second. This moderate clock rate permits the logic within the controller 20 to be implemented with off-the-shelf CMOS EPLD's or similar devices, thus making the cost of the controller 20 (in large quantities) very reasonable.

The movement of request vectors from the input interfaces 12₀-12₂₅₅ to the controller 20 and the movement of connect vectors from the controller 20 to the crossbar switches 15₀-15₆₃ of the switch fabric 14A is a challenging task, because large amounts of control information must be transported every 176 nano second ATM cell interval. For example, in an ATM switch containing 256 input interfaces, 256 16-bit request vectors must be transported to the controller 20 every 176 nano second, leading to an aggregate bandwidth of 23 Gigabits per second between the input interfaces sub-system and the controller 20 sub-system. In addition, 1024 16-bit connect vectors must be transported to the switch fabric 14A every 176 nano second to control the crossbars switches 15₀-15₆₃. This requires an aggregate bandwidth of 93 Gigabits per second between the controller 20 subsystem and the switch fabric 14A sub-system. This 93 Gigabits per second connect vector information can be compressed into 29 Gigabits per second (given that only one input can be routed to an output during each ATM cell interval) by standard compression techniques. However, since this control information should be delivered with high reliability, all of the control connections or control links

between these sub-systems should be dually redundant (not shown in FIG. 4), so there is actually 46 Gigabits per second of data moving between the input interfaces cards and the controller 20 and 58 Gigabits per second of data moving between the controller 20 and the switch fabric 14A. Preferably, high-speed serial links 22 will be used to transmit this control information. For such a case, input interfaces 12₀-12₂₅₅ would be grouped by fours such that only sixty-four serial links would be required to move request vectors from the input interfaces 12₀-12₂₅₅ to the controller 20, and 128 serial links would be required to move the resulting connect vectors from the controller 20 to the pipes 18₀-18₃ (assuming the aforementioned data compression techniques are applied to the connect vectors).

While the use of out-of-band control techniques does require the additional hardware cost of these high-speed serial control links 22, these links 22 cause very little increase the overall system hardware cost. Considering that the 256-input ATM switch 10A of FIGS. 4 and 5 already has 1024 high-speed serial links required to route ATM cells between the input interfaces 12₀-12₂₅₅, and the switch fabric 14A (when the fanout of four is included) and 1024 more high-speed serial links are used to route ATM cells from the switch fabric outputs 19₀-19₁₀₂₃ to the output packet modules 16₀-16₁₅. Thus, the addition of the 192 serial links 22 for routing of the control information increases the total number of high-speed serial links within the system by merely nine percent.

Applying the calculations of Yeh et al. from the article "The Knockout Switch" the ATM cell loss probability of the ATM switch 10A shown in FIGS. 4 and 5 is 4.34×10^{-3} , assuming that the connections of the inputs is symmetrical and not independent as set forth in our co-pending application entitled "TERABIT PER SECOND DISTRIBUTION NETWORK". This cell loss probability falls short of the acceptable ATM cell loss probability of less than 1×10^{-12} mentioned previously.

To reduce the ATM cell loss probabilities, controller 20 applies a temporal spreading technique known as rolling, which provides many statistical advantages. Rolling involves and fulfills three fundamental goals that are aimed at providing more evenly distributed traffic loads. These goals are: (1) spatially distribute the traffic evenly across all pipes 18₀-18₃ so that one pipe will only carry its proportional fraction of the traffic load, (2) spatially distribute the traffic evenly across all of the 16×16 crossbar switches 15₀-15₆₃ within each pipe 18₀-18₃ so that each of the crossbar switches is equally loaded, and (3) temporally distribute the traffic that arrives in a given ATM cell period across two ATM cell periods so that the traffic load can be effectively decreased in an occasional ATM cell period when an unusually high volume of traffic exists and is destined for a particular output packet module. This effective lowering of the traffic load is accomplished by delaying some of the ATM cells arriving during a congested ATM cell interval. The cells are delayed until the next consecutive ATM cell interval when the traffic load competing for the popular resources, i.e. connections to popular output packet modules, will most likely be lower, so the delayed cells should have a higher probability of being routed in the next ATM cell interval. Since the switch fabric 14A is memoryless, the ATM cells that must wait for the next ATM cell interval are stored in their respective input interfaces 12₀-12₂₅₅.

In addition to satisfying these three fundamental goals of packet traffic control to distribute the load, rolling also satisfies two further very important ATM system goals. First, goal (4) is that the ATM switch 10A must guarantee that

ATM cell ordering can be simply maintained when an ATM stream is re-constructed at an output packet module 16₀-16₁₅ even if rolling causes some of the ATM cells within the stream to be delayed differently than others. Secondly, goal (5) is that rolling must also guarantee that the controller 20 will attempt to route every ATM cell through each of the four paths to its desired output packet module, but each of the successive path hunt attempts must occur in a more lightly-loaded 16x16 crossbar switch so that the first attempt occurs in a 16x16 crossbar switch with many previously-routed ATM cells (and very few available paths to output packet modules) while the fourth and final path hunt attempt occurs in an 16x16 crossbar switch that is virtually empty (thereby providing many available paths to output packet modules). The rolling technique is similar to spatial path hunt techniques that pack as many calls as possible in one portion of a spatial network, which by forcing near 100% occupancy in parts of a system results in the remainder of the calls having a very high probability of being successfully routed through the remainder of the system if usage is below 100%. Thus, rolling in its fourth and final path hunt attempt provides a very high probability of an ATM cell successfully being routed. Goal (5), by packing many ATM cells in one portion of the network, superficially seems to conflict with goal (1) that requires the traffic be spatially distributed across the network. However, as will be explained below, temporal spreading provided by the rolling technique permits the network to simultaneously satisfy both goals (1) and (5).

Assuming that each of the 256 input ports 17₀-17_{N-1} of FIG. 4 has an ATM cell that needs to be routed through the distribution network, and assuming that the switch fabric 14A is composed of four pipes 18₀-18₃, then the out-of-band controller 20 may be required to perform 256x4=1024 unique path hunts for the ATM cells before the cells can be routed. To distribute the ATM cells evenly across all four pipes, the 256 ATM cells requesting connections, the rolling technique divides the requests into four groups of equal size. The first group will have path hunts performed for its ATM cells in pipe 18₀ first, then in pipe 18₁, then in pipe 18₂, and finally in pipe 18₃. The second group will have path hunts performed for its ATM cells in pipe 18₁ first, then in pipe 18₂, then in pipe 18₃, and finally in pipe 18₀. The third group will have path hunts performed for its ATM cells in pipe 18₂ first, then in pipe 18₃, then in pipe 18₀, and finally in pipe 18₁. The fourth group will have path hunts performed for its ATM cells in pipe 18₃ first, then in pipe 18₀, then in pipe 18₁, and finally in pipe 18₂. This ring-like ordering of the path hunts guarantees that the routed ATM cells are distributed evenly across all four pipes. In addition, if the ATM cells within each of the four equally sized groups are selected such that the ATM cells within a single group can be routed into exactly four of the 16 inputs on any 16x16 crossbar switch, then the routed ATM cells will also be evenly distributed across all of the 16x16 crossbar switches.

Referring now to FIGS. 5 and 6, a timing diagram for a rolling technique according to the present invention is described. To satisfy goals (1), (2), and (5) simultaneously, the out-of-band controller 20 uses the time delay/time distribution described in goal (3), and these ATM cell delays required by goal (3) must be provided during each ATM cell interval. In all cases, when a group of ATM cells is passed around the ring-like structure of controller 20 from pipe 18₃ to pipe 18₀, the controller 20 re-assigns the cells to the next ATM cell interval (period) which requires that the ATM cells be delayed by one cell period. Because of this re-assignment and delay, each cell group encounters a very lightly-loaded set of 16x16 crossbar switches for its fourth and final path hunt. An additional advantage of this rolling technique using

re-assignment and delay of ATM cell intervals is that it also allows more than 64 simultaneously arriving ATM cells to be routed through the switch fabric 14A to any single output packet module 16₀-16₁₅ (even though there are only 64 connections or links from the switch fabric 14A to each output packet module 16₀-16₁₅). This is occurs with the rolling technique because all of the ATM cells do not need to be routed during the same ATM cell interval. Thus, the rolling technique when used in the out-of-band controller 20 results in extremely low cell loss probabilities both within the switch fabric 14A and the output modules 16₀-16₁₅, even during a transient cell interval that has an extraordinarily high traffic load.

The one ATM cell period delays incurred by some of the ATM cells as they are routed through the switch fabric 14A would normally lead to the conclusion that there would be difficulties in satisfying goal(4) of maintaining proper cell ordering. However, the ring-like ordering of the path hunts within the out-of-band controller 20 guarantees that delayed cells in a stream of ATM cells will always be routed through lower-numbered pipes than non-delayed cells (where pipe 18₀ is the lowest-numbered pipe and pipe 18₃ is the highest-numbered pipe). This information, coupled with the fact that ATM cells are delayed by at most one cell period, ensures that proper cell ordering will be maintained if the cells are extracted from the switch fabric 14A and loaded into first-in-first-out queues 174₀-174₆₃ (shown in FIG. 7) of each output module of the output modules 16₀-16₁₅ in the order of the lowest numbered pipe to the highest numbered pipe: pipe 18₀, pipe 18₁, pipe 18₂, and pipe 18₃.

Referring now to FIG. 7, the output module 16₀ (and the fifteen other output modules 16₁-16₁₅) may be a 64x16 embodiment of the concentrator described in U.S. patent application Ser. No. 08/242,217, entitled "ASYNCHRONOUS TRANSFER MODE SWITCH ARCHITECTURE", filed May 13, 1994, by Cyr et al. now U.S. Pat. No. 5,412,646 and commonly assigned to the assignee of the present invention, which application is hereby incorporated by reference. The output module 16₀ in FIG. 7 is a specific case of the generalized concentrator shown in FIG. 4 of the above-referenced patent application of Cyr et al. Since the output modules 16₀-16₁₅ are well described in the above referenced application, in the interest of brevity they will not be further described here.

To provide a better understanding the equation of the rolling technique, a real-life analogy will be described with respect to FIG. 8, which is a plan view of an amusement park system 500. Consider the problem of transporting a large number of people from amusement park parking lots 511, 512, 513, or 514 to the amusement park 520 using trams to shuttle the people between the two points. Tram system 530 is composed of four tram shuttle trains each with a predetermined route, which is analogous to the four pipes of switch fabric 14A. Each tram shuttle train contains sixteen cars (representing the 16x16 crossbar switches within a particular pipe), and each shuttle car is equipped with sixteen seats (representing the output links emanating from a single 16x16 crossbar switch). In this analogy, each customer (representing an ATM cell) arrives in one of four parking lots 511, 512, 513, or 514 surrounding the amusement park 520. As a result, each customer is instantly placed in one of four groups, and since the parking lots 511-514 are the same size, each group contains an equal number of customers on the average. The customers in any single parking lot 511, 512, 513, or 514 must then divide up and stand in one of sixteen lines, where each line is associated with a respective car of the tram shuttle train. The amusement park 520 is sub-divided into sixteen different theme areas (The Past Land, The Future Land, etc.), and each of the sixteen seats of a particular tram car is labeled with the

theme area to which that seat's occupant will be given admission. Before arriving in the parking lot, each customer must randomly chose one of the sixteen theme areas (representing the sixteen output packet modules 16₀-16₁₅) where he or she wishes to spend the day. Customers must then find an available seat associated with their desired theme area on one of the four trams that passes by the loading area 531, 532, 533, or 534 of their parking lot. If a customer has not found an available seat after four trams have passed by, then he or she is not permitted to enter the amusement park during that day (This harsh condition represents the loss of an ATM cell due to blocking in all four pipes of the distribution network, a small but finite possibility).

The first tram that stops at the loading area that the customer can try has already visited three other parking lot loading areas, so the customer's pre-specified seat may be full. However, if the customer does find his or her seat to be vacant on that tram, then the tram will deliver him or her straight to the amusement park 520. If the customer fails to get on the first tram, he or she must wait and try the second tram which has already visited two other parking lot loading areas. If the customer is successful at finding his or her pre-specified seat on the second tram, that tram will deliver the customer to the amusement park 520 after one more parking lot stop. If the customer fails to get on the first tram and the second tram, then he or she must wait and try the third tram which has only visited one other parking lot loading area. If the customer is successful at finding his or her seat on the third tram, that tram will deliver him or her to the amusement park 520 after two additional parking lot stops. If the customer fails to get on any of the first three trams, then the customer must wait and try the fourth and final tram. Fortunately, this tram has not visited any parking lots yet, so the arriving tram is empty, and the customer's seat will be taken only if another customer in his/her parking lot line is also trying for the same seat. The system 530 satisfies goal (5), because each of the successively arriving trams is more lightly-loaded than the previous one. Thus, a controller 20 rolling ATM cells indeed can fulfill goals (1), (2), and (5).

The rolling technique if used by itself improves the ATM cell loss probability of ATM switch 10A from 4.34×10^{-3} to approximately 10^{-11} . Using the analysis techniques of the article "A Growable Packet Switch Architecture" the cell loss probabilities for an ATM switch 10A that has independent connections to the inputs of the switch fabric 14 according to Galois field theory and also has an out-of-band controller 20 that incorporates rolling techniques can be analytically modeled and calculated. Each of the 16x16 crossbar switches in pipe 18₀ receives an offered traffic load equal to $R_a = R_L/4 + R_{res}$, where R_{res} is defined to be the fraction of the 16 inputs to a 16x16 crossbar switch that are blocked in pipe 18₁ and routed to pipe 18₀ for a re-attempt. For a first attempt at solving for the cell loss probability, let us assume that $R_{res} = R_L/16$. Thus, the cell loss probability of a single 16x16 crossbar switch in pipe 18₀ can be determined using the equation of Eng et al.

$$P(\text{cell loss}) = \left\{ 1 - \frac{m}{(nR_L)} \left[1 - \sum_{k=0}^m \frac{(nR_L)^k e^{-nR_L}}{k!} \right] \right\} + \frac{(nR_L)^m e^{-nR_L}}{m!};$$

where $m=1$, $n=1$, and the switch loading is given by $R_a = R_L/4 + R_L/16$. Using these assignments, the resulting cell loss probability for a fully-loaded ($R_L=1.0$) pipe 18₀ 16x16 crossbar switch can be calculated to be:

$$P(\text{cell loss in pipe } 18_0) = 1.3 \times 10^{-1}.$$

Thus, the fraction of the 16 inputs to a 16x16 crossbar that are passed to the second pipe after the first attempt is given by:

$$\beta - 2 = R_a \times P(\text{cell loss in pipe } 18_0) = (3.13 \times 10^{-1})(1.3 \times 10^{-1}) = 4.06 \times 10^{-2}.$$

By symmetry, this should have also been the same as the fraction of inputs that are passed from pipe 18₁ to pipe 18₀, so the residue assumption of $R_L/16=0.062$ above was incorrect. By refining this assumption and performing a second attempt, and now assuming that $R_{res} = R_L/32$. Thus, the cell loss probability of a single 16x16 crossbar switch in pipe 18₀ can be determined again using the equation of Eng et al., where $m=1$, $n=1$, and the switch loading is given by $R_a = R_L/4 + R_L/32$. Using these assignments, the resulting cell loss probability for a fully-loaded ($R_L=1.0$) pipe 18₀ 16x16 crossbar switch is calculated to be:

$$P(\text{cell loss in pipe } 18_0) = 1.2 \times 10^{-1}.$$

Thus, the fraction of the 16 inputs to a 16x16 crossbar that are passed to the second pipe after the first attempt is given by:

$$\beta - 2 = R_a \times P(\text{cell loss in pipe } 18_0) = (2.81 \times 10^{-1})(1.2 \times 10^{-1}) = 3.37 \times 10^{-2}.$$

This calculation result is very, close to the assumed value of $R_{res} = R_L/32 = 3.13 \times 10^{-2}$, so the assumption is considered to be satisfactory. The blocked cells are sent to pipe 18₁ for subsequent path hunting, and they encounter a negligible number of ATM cells from previous attempts. Thus, the 16x16 crossbar switch in pipe 18₁ can be modeled for analysis as a growable packet switch, with $m=1$, $n=1$, and $R_a = \beta - 2$, and the resulting cell loss probability of this model is 1.4×10^{-21} . The fraction of the 16 inputs to the 16x16 crossbar in pipe 18₁ that are passed to the pipe 18₂ is 4.2×10^{-4} . Similar arguments can be used to show that the resulting cell loss probability for cells entering pipe 18₂ is 1.9×10^{-4} , and the resulting fraction of the 16 inputs to a 16x16 crossbar passed to pipe 18₃ is 7.9×10^{-8} . The resulting ATM cell loss probability in pipe 18₃ is 3.7×10^{-8} , and the fraction of the 16 inputs to a 16x16 crossbar not routed in pipe 18₃ (and therefore not routed in all four pipe attempts) is 2.9×10^{-15} . Thus, through the use of the rolling techniques within the out-of-band controller 20, the ATM cell loss probability of an ATM switch 10A with independent connections at the inputs of its switch fabric 14A can be decreased from an unacceptable value of 1.47×10^{-6} to an acceptable value of 2.9×10^{-15} .

A preference technique may be used in conjunction with the rolling technique described above to decrease the cell loss probability of an ATM switch 10A even further. Referring back to FIG. 8 and the amusement park analogy, some form of arbitration was required at the tram loading areas to determine which of the customers in the line will be given a particular seat on the tram when more than one customer is requesting the same seat. Similarly, the out-of-band controller 20 must provide an arbitration scheme for selecting which of the arriving ATM cells will be assigned a particular link whenever two or more cells request access to the same link. The arbitration scheme used can have an advantageous effect on the ATM cell loss probabilities.

One possible arbitration scheme is a random scheme to determine which of the ATM cells is assigned the link. The random selection scheme is the scheme assumed for the analysis of the rolling technique presented above. However, other arbitration schemes are possible, and one particular

arbitration scheme that has advantageous results is called the preference scheme. The preference arbitration scheme assigns a preference weight to each of the ATM cells in a particular grouping. ATM cells with higher preference weights are given precedence over ATM cells with lower preference weights whenever two or more cells request access to the same link. As a result, an effective hierarchy is created within the groupings of ATM cells.

The creation of a hierarchy may superficially seem to produce undesirable characteristics within the switch fabric 14A, because customers with high preference weights will be offered better service than customers with low preference weights. In fact, the one customer with the highest preference weight within each group can never have his or her ATM cell blocked by another customer's ATM cell. Although this may seem unfair, a detailed analysis of the effects of imposing this hierarchy indicates that it actually leads to improved performance, i.e. lower cell loss probabilities, for all customers—even for the customer at the bottom of the hierarchy with the very lowest preference weight.

The results of this analysis are summarized in FIG. 9, where the probability of loss of an ATM cell; i.e., the probability of a cell not being assigned to an available path, is shown as a function of the number of path hunts that were attempted in different pipes by the out-of-band controller 20. In this analysis, it was assumed that the group sizes were four—i.e., up to four ATM cells could simultaneously compete for access to the same link. As a result, four different preference weights were assigned to create a hierarchy for the four input ports associated with each group. The preference weight associated with a particular input port is assumed to be a fixed constant that does not vary with time. The resulting plots 901, 902, 903 and 904 in FIG. 9 indicate that the cell loss probability decreases as more path hunts in more pipes are performed, but it also shows that the inputs with the lower preference weights 903, 904 have higher cell loss probabilities than the inputs with higher preference weights 901, 902, as might be expected. Super-imposed on these plots is a similar plot 910 which indicates the probability of not being served when a random selection arbitration scheme is used instead of the hierarchy arbitration scheme. The surprising and unexpected results are that after path hunt attempts in four different pipes, the random selection arbitration scheme produces cell loss probabilities which are higher than the average of the cell loss probabilities for the hierarchy arbitration scheme. In fact, the plot 910 of the random selection arbitration scheme shows an average cell loss probabilities for all of the input ports which are notably higher than the plots 903 and 904 which are the average cell loss probabilities for even the input ports with the lowest preference weights within the hierarchy arbitration scheme. This phenomenon can be explained by the fact that after three sets of path hunts in three different pipes, the distribution of ATM cell requests entering the fourth pipe is very different depending on whether the random or preferences arbitration scheme is used. In the random selection arbitration scheme, there is a small but equal probability that all of the ATM cells are requesting a path. However, in the hierarchy arbitration scheme, most of the ATM cells with higher preference weights will be requesting a path with a probability of practically zero, while the ATM cell with the lowest preference weight will be requesting a path with a sizable probability, because that particular ATM cell may have been denied access to links in all three of its previous path hunt attempts. However, a single request arriving with a high probability at the fourth and last path hunter in the

controller will lead to more routed ATM cells than many requests arriving with low probability, because the single request can always be satisfied since contention for an output link will never occur.

As a result, it seems apparent from the plots in FIG. 9 that by assigning preference weights to the input ports and by using a hierarchy arbitration method to resolve link contention and route paths in the out-of-band controller, the worst-case cell loss probability of the switch fabric 14A can be decreased from 2.9×10^{-15} that was achieved by the introduction of the rolling technique to an even lower value of 2.4×10^{-16} . It is worth noting that input ports that are assigned higher preference weights will encounter even lower cell loss probabilities as indicated in FIG. 9.

Referring back to FIG. 5, in order to provide a physical embodiment of the rolling and preference methods, the ATM switch 10A is segmented into four basic sub-systems. These four sub-groups consist of the input interfaces 12₀–12₂₅₅, the output modules 16₀–16₁₅, the switch fabric 14A, and the out-of-band controller 20.

The input interfaces 12₀–12₂₅₅ within the network provide the necessary interfaces between the incoming transmission links and the links connected to the switch fabric 14A and the out-of-band controller 20. As a result, the input interfaces 12₀–12₂₅₅ must provide a termination for the input transmission line. For example, if the input transmission line is a SONET link, then the input interface must provide for clock recovery, link error detection, SONET pointer processing and frame delineation, ATM cell extraction, and an elastic storage function to synchronize the arriving ATM cells to the system clock within the distribution network. The extracted ATM cells are then loaded into a FIFO buffer of the input interface. The input interface must also read ATM cells from the FIFO buffer and extract the ATM header from the cell. The VPI/VCI field of each ATM header is then used as an address into a translation table located on the input interface. The output of the translation table provides a new VPI/VCI field and the address of the output packet module to which the ATM cell is to be routed. The new VPI/VCI field is written into the ATM cell as a replacement for the old VPI/VCI field, while the output module address is routed as a request vector to the out-of-band controller 20 for the controller fabric 14A. Since the amount of processing time required by the out-of-band controller 20 is a fixed value, the input interface simply holds the ATM cell in a buffer until the out-of-band controller 20 has completed its path hunt and has relayed the results into the switch fabric 14A. Once the switch fabric 14A is loaded with the new switch settings to appropriately route the ATM cell, the input interface can inject the ATM cell into the switch fabric 14A and it will be automatically routed through the switch fabric 14A to its desired output module 16₀–16₁₅. It should be noted that each input interface 12₀–12₂₅₅ actually is provided with one link to each of the four pipes 18₀–18₃ of the switch fabric 14A. In addition, the use of rolling (i.e. temporal spreading) within the switch fabric 14A may require a copy of the ATM cell to be injected into each of the four links during any one of two consecutive ATM cell intervals. As a result, the timing within the input interfaces 12₀–12₂₅₅ must be tightly coupled and synchronized to the timing of the rest of the sub-systems within the ATM switch 10A.

Each of the two hundred fifty six input interfaces 12₀–12₂₅₅ in FIG. 5 are numbered with an address ranging from 0 to 255, but each input interface is also assigned an alias address given by a letter between A and P. These alias addresses are used to identify which input port the input

15

interfaces will connect to within the switch fabric 14A. The actual set of four crossbar switches to which a particular input interface is connected is determined by the Galois field techniques that were described previously. These techniques guarantee independence between all of the inputs on any 16x16 crossbar switch of any pipe.

Each of the sixteen output modules 16₀-16₁₅ in FIG. 5 is labeled with addresses ranging from AA to PP, and each output module performs an important function within the ATM switch 10A. Each of the output modules 16₀-16₁₅ within FIG. 5 provides terminations for a respective set of sixty-four links emanating from the switch fabric 14A. Each output module 16₀-16₁₅ also provides two basic functions: it provides a small degree of space switching to route each ATM cell arriving on one of the sixty-four inputs to the desired one of the sixteen output ports, and it provides buffering of ATM cells to handle the problems associated with multiple packets that are simultaneously destined for the same output Out₀-Out₂₅₅.

There are many ways for these two functions to be implemented. The most straight-forward approach would probably construct a shared memory switch that could perform sixty-four memory writes and sixteen memory reads within an ATM cell interval (176 nano second). The memory could then be treated as sixteen disjoint linked lists (one for each output Out₀-Out₂₅₅) along with a seventeenth linked list containing idle memory locations. Although simple, this approach requires eighty memory accesses every 176 nano second, so it would demand memories with 2.2 nano second access times. An alternate approach would split each 64x16 output module 16₀-16₁₅ into a 64x16 concentrator and a 16x16 shared memory switch. The concentrator would be a memory system that provides for sixty-four writes and sixteen reads every ATM cell interval, but the memory size could be small (and memory speeds could be fast) since the buffering required for output contention problems is not provided in this memory. In addition, the 64x16 concentrator could be implemented as a single linked list spread out across sixty-four distinct memory chips. As a result, each memory chip would require only one write and up to sixteen reads for every ATM cell interval. The 16x16 shared memory switch only performs thirty-two memory accesses every ATM cell interval, so slower (and larger) memories could be used, and the buffering for output contention problems could be provided in this shared memory portion of the output module. Thus, this latter arrangement is the more practical alternative for an output module.

The switch fabric 14A is essentially a group of small circuit switches that provide the required connectivity between the input interfaces and the output modules in response to the control signals generated by the out-of-band controller 20. In the embodiment of the ATM switch 10A shown in FIG. 5, the switch fabric 14A is composed of sixty-four 16x16 crossbar switches, where disjoint groups of sixteen switches comprise a pipe. The four pipes are labeled pipe 18₀, pipe 18₁, pipe 18₂, and pipe 18₃, and the sixteen 16x16 crossbar switches within a given pipe are labeled switch 0-15. The crossbar switches must be capable of receiving the control signals generated by the out-of-band controller 20 and must reconfigure all of the switch settings during a guard-band interval between consecutive ATM cells. Each 16x16 crossbar switch supports sixteen inputs labeled input A through input P, and each 16x16 crossbar switch also supports sixteen outputs labeled output AA to output PP. It was noted above that each input interface connects to a different 16x16 crossbar in each of the four

16

pipes 18₀-18₃, but it should now be noted that an input interface that connects to input X in pipe 18₀ is required to be connected to input X in the other three pipes 18₁-18₃, as well, where X is an element of the set {A,B,...,P}. The actual connections between the input interfaces 12₀-12₂₅₅ and the crossbar switches within the switch fabric 14A are determined using Galois field theory techniques that were referenced above. These techniques guarantee independence between input ports for routing within switches in each pipe of the switch fabric 14A. FIG. 5 also illustrates that output YY from each of the sixty-four crossbar switches is routed to one of the sixty-four inputs on the 64x16 output module labeled YY, where YY is an element of the set {AA,BB,...,PP}.

The basic function of the out-of-band controller 20 for the switch fabric 14A is to determine through which of the four pipes 18₀-18₃ a particular ATM cell may be routed. Once the out-of-band controller 20 has successfully determined a pipe through which the ATM cell is to be routed without being blocked, the task of setting up the path through the pipe is simple, because by the definition of a pipe, there will exist only one path within the pipe between the input port of the arriving ATM cell and the desired output module. As a result, the fundamental path hunting task of a switching network is essentially reduced to the simpler task of pipe hunting in the ATM switch 10A.

The out-of-band controller 20 still requires a large busy-idle table to identify the status of each of the intermediate (FN) links between the 16x16 crossbar switches of the switch fabric 14A and the output modules 16₀-16₁₅ as busy and unavailable or idle and available. However, this large busy-idle table may be sub-divided into many small busy-idle tables that the controller 20 can access in parallel, and thereby perform many pipe hunting operations in parallel. There are many ways to implement the controller 20 for a large switch having the general growable packet switch architecture. In the extreme case, four levels of parallelism may be applied to the architecture of the controller 20 to perform pipe hunting. One embodiment that uses three levels of parallelism will be described in detail, first and then a fourth level of parallelism for the controller 20 will be discussed.

The first level of parallelism is obtained by providing each of the four pipes 18₀-18₃ with a respective pipe hunt controller 24₀-24₃. This level of parallelism allows pipe hunting to be carried out in all four pipe hunt controllers 24₀-24₃ simultaneously. The second level of parallelism is obtained by providing switch controllers 26₀-26₆₃, with sixteen switch controllers within each pipe hunt controller 24₀-24₃. A unique switch controller 26₀-26₆₃ is respectively associated with each of the 16x16 switches within each pipe of the switch fabric 14A. As a result, pipe hunting operations can be carried out in parallel within all sixteen of the switch controllers of each pipe hunt controller 24₀-24₃. The third level of parallelism is obtained by permitting each of the switch controllers 26₀-26₆₃ to perform parallel processing over all sixteen of the output links attached to its respective 16x16 crossbar switch. Effectively, each of the switch controllers 26₀-26₆₃ reads sixteen busy-idle bits from its busy-idle memory in parallel, performs parallel pipe hunting operations based on those sixteen bits, and then writes the sixteen resulting busy-idle bits into its respective busy-idle memory in parallel with the other busy-idle memories. A representative switch controller 26₀ of the sixty four switch controllers 26₀-26₆₃ is shown in FIG. 10. The concurrent processing of sixteen busy-idle bits is accomplished by providing switch controller 26₀ sixteen unique link control-

lers AA-PP, each of the link controllers AA-PP is assigned the task of processing busy-idle bits for one intermediate link between its portion of the switch fabric 14A and its respective output modules. In the embodiment shown in FIG. 10, the large busy-idle memory required to control 5 switch 10A has been divided into many single bit memories, busy-idle flip-flops, with each single bit, busy-idle memory being logically and physically associated with its respective link controller AA-PP.

The general data flow for request vectors generated by the input interfaces 12₀-12₂₅₅ is shown in FIG. 5. For example, input interface 12₀ in FIG. 5 routes its request vector via connection 21₀ to pipe hunting controller 24₀ where it is poked into the pipe hunting ring (i.e. controller 20), and the rolling scheme requires the request vector to be looped 15 through pipe hunt controller 24₁, pipe hunt controller 24₂, and pipe hunt controller 24₃ as it circulates around the ring. In general, each of the input interfaces 12₀-12₂₅₅ produces one request vector, and each request vector will contain a number of bits equal to the number of output modules within the system. The request vector from a single input interface in FIG. 5 is thus a sixteen-bit data word, where each bit of the request vector points to one of the sixteen output modules. If an ATM cell within a input interface is requesting a connection to an output port on the i-th output module, then bit i within the request vector will be set to a logic "1" 25 and all other bits within the request vector will be set to a logic "0". When the controller 20 receives this particular request vector from the input interface, it can then identify that a path is required between the source input interface and the i-th output module. 30

The entire sixteen-bit request vector from a input interface is routed via a respective control connection 21₀-21₂₅₅ to one of the four pipe hunt controllers 24₀-24₃, and the controller 20 pokes the vector into one of the sixteen switch 35 controllers associated with that particular pipe hunt controller. As shown in FIG. 10, the sixteen bits of the request vector are injected into a switch controller and are distributed across all sixteen of the link controllers within that particular switch controller. Each link controller is associated with a single link between the crossbar switches and the output modules, and it essentially processes one bit of the sixteen-bit request vector. This finite state machine circuitry that is associated with a single link controller consists of one flip-flop (the single-bit memory required to store the busy-idle bit associated with this link controller's link) and four 45 logic gates. A state table description of the link controller operation is given in FIG. 12, where the state variable is defined by the busy-idle bit. The link controller hardware provides for one request vector input bit, designated request-in; one request vector output bit, designated request-out; and one connection vector output bit, designated connect. The request vector input bit is a logic "1" if the input desires a connection through the link associated with this link controller-otherwise, it is a logic "0". The request vector output 55 bit is a logic "1" if the logic "1" input request vector bit was not satisfied by this particular link controller-otherwise, it is a logic "0". The connect vector output bit is a logic "1" if the logic "1" input request vector bit was satisfied by this particular link controller indicating the ATM cell will be routed to its desired output module through the link associated with this link controller-otherwise, it is a logic "0". The busy-idle flip-flop in FIG. 10 is reset to the logic "0" (idle) state at the beginning of each ATM cell slot, so the first request vector bit that enters the link controller with a logic 65 "1" request is assigned the link (creating a logic "1" connect vector bit and a logic "0" output request vector bit) and sets

the busy-idle flip-flop to the logic "1" (busy) state. Any subsequent request vector bits that enter the link controller during this particular ATM cell slot will be denied a connection through this link (forcing a logic "0" output on the connect vector bit and creating an output request vector bit that is identical to the input request vector bit). A time-lapsed view of several consecutive sixteen-bit request vectors passing through a single switch controller is shown in FIG. 12, along with the resulting states of the busy-idle bits stored within the switch controller. The resulting output request vectors and output connect vectors illustrate the general operation of each of the pipe hunt controllers 24₀-24₃.

The use of rolling within the controller 20 requires a very precise temporal ordering of two fundamental events: poking and busy-idle flip-flop clearing. The timing diagram of FIG. 13 illustrates the synchronization and data flow that might be used for the logic within the controller 20. As indicated by the timing diagram, the flow of data around the ring of controller 20 is from pipe controller 24₀ to pipe controller 24₁, to pipe controller 24₂, to pipe controller 24₃, and back to pipe controller 24₀. Request vectors generated by input interfaces with alias addresses A, B, C, and D are poked into pipe controller 24₀. Request vectors generated by input interfaces with alias addresses E, F, G, and H are poked into pipe controller 24₁. Request vectors generated by input interfaces with alias addresses I, J, K, and L are poked into pipe controller 24₂. Request vectors generated by input interfaces with alias addresses M, N, O, and P are poked into pipe controller 24₃. The poking times and busy-idle bit clearing times take place at different moments within each of the pipe hunt controllers 24₀-24₃. From the point of view of any pipe controller, the request vector bits flow through the pipe controller in alphabetical order (A to P) if one ignores the busy-idle bit clearing times. This ordering guarantees that the aforementioned advantages of preferences will be realized within the controller 20, because the request vector generated from a input interface with alias address A will always be given precedence over the request vectors generated from input interfaces with alias addresses B, C, and D, etc.

The benefits derived from forced independence between the inputs on a particular 16x16 crossbar switch produce a slight increase in the complexity of the pipe hunter circuitry. Because of the independent connections between the input interfaces and the switch fabric 14A, which independence is assured by the use of Galois field theory, a request vector from a single input interface must be appropriately routed to several different switch controllers in each of the stages in the pipe hunting ring. The mixing nature of the Galois field theory generated connections requires each input interface 12₀-12₂₅₅ to be connected to a different set of 16x16 crossbar switches within the switch fabric 14A, and as a consequence, it also requires request vectors generated on different input interfaces to be routed through entirely different sets of switch controllers within the controller 20. Since request vectors are time-multiplexed on links within the controller 20, all of the request vectors (within a particular ATM cell slot) that are expelled from a particular switch controller in one pipe hunter stage must (by definition) be routed to different switch controllers in the next pipe hunter stage. To provide this dynamic routing of the request vectors, each pipe hunt controller 24₀, 24₁, 24₂ and 24₃ is connected to a respective small switching network 30₀, 30₁, 30₂ and 30₃, shown in FIG. 5. Alternatively, simple multiplexers may be used instead of switching networks 30₀, 30₁, 30₂ and 30₃, thereby greatly decreasing costs for the controller 20. Fortunately, the required configurations of these

small switching networks 30₀, 30₁, 30₂ and 30₃, (or multiplexers) are cyclic with a period equal to the ATM cell period, and the required configurations can be determined a priori and can therefore be "hard-coded" into the small switching networks (multiplexers) during the design of the circuitry of the controller 20.

As mentioned previously, ATM switch 10A shown in FIG. 5 might be scaled such that the number of input lines were 512, 1024 or even higher. For those size switches, assuming that the input lines are carrying 2.5 gigabits per second data rates, the aggregate throughput would be over 1.0 terabits per second. For switches of that size, a fourth level of parallelism may be needed to provide sufficient processing power for the controller 20 to hunt for all those paths through all those pipes. For ATM switches with 512 and 1024 input lines, the data rates on wires within their respective controllers are 204 megabits per second and 386 megabits per second, which is considerably higher than the 113 megabits per second rate of the 256 input line version of ATM switch 10A.

The basic idea behind the fourth level of parallelism is a modification of the previously described controller 20 design which requires that request vectors be routed through the pipe hunter stages in parallel. In particular, all of the request vectors that are poked into a particular pipe are routed through the pipe hunter stages together, and these request vectors are said to comprise a poke group. In the embodiment shown in FIG. 5, this approach to the design of controller 20 creates four poke groups of sixteen-bit request vectors, so each poke group contains sixty-four bits. The four poke groups can be labeled with a concatenation of the four alias labels on the request vectors. As a result, the four poke groups for the re-designed pipe hunter of FIG. 5 are called ABCD, EFGH, IJKL, and MNOP. It is important to note that whenever a single sixty-four bit ABCD poke group is being routed through one of the switch controllers in pipe controller a of FIG. 5, there is also a sixty-four bit ABCD poke group being routed through each of the other fifteen switch controllers in pipe controller 24₀. As a result, there are a total of 1024 request vector bits associated with sixteen ABCD poke groups that are being routed through pipe 18₀ at a single instant of time. The modified controller 20 processes the request vectors for all N input ports (by passing them through all four pipe hunt controllers 24₀–24₃) every eight clock cycles, and since this task must be completed within a single 176 nano second ATM cell interval, the required clock rate within the controller 20 is 46 megabits per second regardless of the size (aggregate throughput) of the N×N ATM switch. As a result, since the controller 20 must perform eight processing steps (regardless of the network size), the process is said to be an O(1) path hunt algorithm. During the execution of this O(1) path hunt algorithm for the N=256 input ATM switch 10A of FIG. 5, the equivalent of 16,384 link controller path hunts and 16,384 link controller path hunt checks are performed every 176 nano second, so if each path hunt is considered to be an instruction execution and each path hunt check is considered to be an instruction execution, then the controller 20 can be viewed as a parallel processor capable of sustaining a 186 giga-instructions per second processing rate. The trade-off for maintaining a reasonable data rate in the controller 20 (regardless of size) is an increase in link controller logic complexity and an increase in signal connections passing between successive stages of the controller 20 as the size is increased. ATM switch designs with aggregate throughputs in excess of 1 Terabits per second will require between 4096 and 32,768 signals at 46 megabits per second to be routed

between successive stages of the controller, i.e. between pipe hunt controllers 24₀–24₃.

In addition to increasing the number of signals between pipe hunt controller stages, the use of parallelism within the controller 20 also requires a slight increase in the hardware requirements for each link controller, because each link controller must now support a parallel path hunt on four bits within the poke group. The extra hardware added to the controller 20 by the fourth level of parallelism should be offset by the resulting lower processing rate.

In addition to large throughputs and low data losses, an important and essential attribute for any switching product that will be used in the public switched network is fault tolerance, in order to provide a very high level of availability. A switching system that is fault-tolerant must display most, if not all, of the following attributes: 1) the ability to detect that a fault exists, 2) the ability to locate and identify the faulty component, 3) the ability to avoid the faulty component by detouring traffic through alternate paths within the network, 4) the ability to provide an acceptable level of performance (cell loss probability) even in the presence of a small percentage of faulty components, and 5) the ability to permit maintenance personnel to easily repair the faulty component (e.g., by swapping a board), and 6) the ability to provide an acceptable level of performance, i.e. cell loss probability, even when the faulty component is being repaired. These attributes typically require some level of redundancy within the switch paths to satisfy the requirements of attributes 3 and 4, and they also require redundancy at a next level higher within the network fabric and controller to satisfy the requirements of attributes 5 and 6.

A large benefit of using out-of-band control techniques within an ATM switch is derived from the fact that existing fault-tolerance techniques that have been applied for years in circuit switches can be re-used in the ATM switch 10A. As a result, all six of the above fault-tolerance requirements can be easily satisfied within the general architecture of ATM switch 10A. For example, FIG. 15 shows a plot of a simulation of the cell loss probability within the architecture of FIG. 5 when faulty links are added to the switch fabric 14A. It can be seen that while the cell loss probability increases as fault links are added, up to 0.5 percent of the links can be faulty before the cell loss probability exceeds the maximum acceptable level of 1×10^{-12} . This is a direct result of the fact that the architecture in FIG. 5 provides four paths between each input port and output port. Faulty paths can be rapidly identified by parity or CRC checks at the output modules 16₀–16₁₅. If an error is detected, the controller 20 already knows what path the corrupted ATM cell had been routed through, so it can check the path by sending an "interrogation ATM cell" through the switch fabric 14A. If the interrogation ATM cell is also corrupted, then the path should be taken out of service by writing the busy-idle bit for the path to the maintenance busy state, which is not cleared even when the global clear is sent at the end of each ATM cell period.

Although the present ATM standard does not define or require a service that supports variable length cells or packets, the evolutionary tendencies of the telecommunication industry are towards ATM cells with lengths that differ from the initially defined 53-byte standard. This change may evolve as a result of the satisfaction (or dissatisfaction) that different users may find as they begin to experiment with different applications being transported over the ATM packet lines and networks. For example, the very fact that the current ATM cell length represents a compromise between the voice and data communities indicates that there may be

some customers in the future who are not entirely satisfied with the offered cell size. Should such customers become organized, it is possible that another cell length (other than 53 bytes) may be requested. For example, the CATV industry is already considering cell sizes greater than 53 bytes for transporting MPEG-2 digital video streams. Although larger packets might be routed inside multiple 53 byte ATM cells, the resulting bandwidth inefficiencies may ultimately lead to the desire for a new cell length standard. Thus, it is desirable to have an ATM switch architecture that can adapt to changes in cell or packet length.

With only a few minor modifications to the input interfaces 12₀-12₂₅₅ and the controller 20, operation of ATM switch 10A with cells of arbitrary cell lengths is provided. This is accomplished by allowing arbitrary length that are an integer multiple of the basic cell period, which can be 53 bytes or some other desired length. The ATM switch 10A may be modified so readily to support arbitrary cell lengths primarily because ATM switch 10A is essentially a circuit switch, i.e. an arbitrarily long message switch, with the update capability of a very fast out-of-band path hunt processor. For fixed cell length operation described previously, the controller 20 of the switch 10A performs path hunts for cells that arrive on potentially all of the N input ports, and it must then set up the N paths for all of these routed cells. At the end of the 176 nano second cell interval, all of these N paths are globally torn down, making all of the network connections idle for the next cell interval when the entire process is repeated. If variable length cells are permitted, then it should be apparent that the global tear-down of all N paths at the end of a cell interval is no longer permitted. In fact, all of the paths from one cell interval must be left established in the next cell interval, and individual path tear-downs must be implemented whenever the termination of a cell is identified by its respective input interface 12₀-12₂₅₅. Thus, to modify ATM switch 10A to handle variable length cells requires that each input interface 12₀-12₂₅₅ be modified to be capable of identifying the start and end of each cell, either with fixed and unique start and end patterns, a unique start pattern coupled with a cell length identifier contained within the cell, or some other type of indication. Each modified input interface must then be capable of sending two different types of request vectors to a modified controller 20': one to request a path set-up and one to request a path tear-down. This can be accomplished by adding a single bit to the 16-bit request vector shown in FIGS. 13A-13D, where the additional bit is used to indicate whether the request is a set-up request or a tear-down request. Upon reception of the request vector, the controller 20' routes the request vector through all four of the pipe controllers 24₀-24₃ (as before), but the link controllers must now be capable of both setting and resetting the busy-idle flip-flops in response to the different types of request vectors. In addition, the link controller associated with a particular output link must also maintain a memory indicating which input presently has a path established to its output link, because only that input is permitted to tear down the path to the output link. The hardware required for all of these functions within a link controller is illustrated in FIG. 16. It should be noted that the inclusion of variable length cell routing within the distribution network requires that the processing rate within the controller 20' be increased by a factor of two, because it is possible that every input may require a single path set-up and a single path tear-down every ATM cell interval. It should also be noted that the inclusion of variable length cell routing does not preclude the implementation of any of the other features of the switch 10A that have been mentioned previously.

The use of variable length cells within the switch fabric 14A also requires that the output modules 16₀-16₁₅ be modified to route cells with different cell lengths. The lengths of the buffers would have to be increased to accommodate at least four times the longest cell or packet that can be communicated.

The rapid acceptance of ATM within the LAN and WAN communities indicates that ATM may already be developing a strong foot-hold in the private-switched network environment. As a result, it may be only a matter of time before there is a strong demand for ATM services in the public-switched network environment. However, when or if that strong demand occurs is uncertain. There is some question within the telecommunications industry of the ability of ATM to efficiently provide services that are inherently constant-bit-rate services (voice and video), and also the ability of ATM networks to effectively route traffic in a network where the traffic is highly correlated rather than random. Because of this uncertainty, ATM service providers and ATM switch vendors must proceed cautiously. A switch based on the switch fabric 14A is a sensible system in such uncertain times because the architecture is flexible enough to provide both packet switching (ATM) and circuit switching (STM) communications at the same time. Such a switch 10B is shown in FIG. 17. Since STM switching is a form of circuit switching, and since the switch fabric 14A is essentially a circuit switch with very fast path hunting capabilities, the switch fabric 14A is well suited for the routing of STM traffic. A slightly different controller 620 is required for STM traffic, and the input interfaces 612 and output modules 616 provide time-slot interchanger functions required by circuit switching equipment, but the single stage switch fabric 14A can remain unaltered in a combination STM and ATM switch, or in a wholly STM switch. Simulations have been written to analyze the operation of the switch fabric 14A in a wholly STM environment, where an N=256 ATM switch is modified to implement an N=128 STM switch. The resulting blocking probability of this N=128 STM switch has been calculated from this simulation to be less than 1×10^{-9} . This is a very acceptable value in a circuit switched environment where packet loss is not an issue. Thus, switch 10B can vary the percentage of STM traffic it carries up to 100 per cent. This flexibility greatly reduces or eliminates the possible financial consequences of the uncertainty in the demands of customers for future ATM and STM services.

The switch fabric 14A is essentially technology-independent. An embodiment using free-space digital optics as the interconnection technology within the switch fabric is contemplated. The 16x16 crossbar switches within the switch fabric 14A will be implemented with FET-SEED device arrays. Such an approach may provide many benefits within the switch fabric 14A, because the resulting design based on optical interconnections may have lower levels of signal crosstalk, lower chip counts due to increased device integration, lower signal skew, and lower overall power dissipation, which results in simpler thermal management techniques within the switch fabric 14A.

While the invention has been particularly illustrated and described with reference to preferred embodiments thereof, it will be understood by those skilled in the art that various changes in form, details, and applications may be made therein. It is accordingly intended that the appended claims shall cover all such changes in form, details and applications which do not depart from the true spirit and scope of the invention.

What is claimed is:

1. A packet switch for switching circuit switched com-

23

munications from a plurality of circuit switched input lines to a plurality of circuit switched output lines and packet switched communications from a plurality of packet switched input lines to a plurality of packet switched output lines, comprising:

- a plurality of circuit switched input interfaces, each having an input port connected to a respective circuit switched input line of said plurality of circuit switched input lines, and each of said circuit switched input interfaces having an output port;
 - a plurality of packet switched input interfaces, each having an input port connected to a respective packet switched input line of said plurality of packet switched input lines, and each of said packet switched input interfaces having an output port;
 - a single stage switch fabric having a plurality of input ports with a first portion of said input ports connected to respective output ports of said circuit switched input interfaces and a second portion of said input ports connected to respective output ports of said packet switched input interfaces;
 - a plurality of circuit switched output modules, said circuit switched output modules together having a plurality of inputs, each of said circuit switched output module inputs connected to respective output port of said first portion of said single stage switching fabric, and together having a plurality of outputs, each of said circuit switched output module outputs connected to a respective circuit switched output line of said plurality of circuit switched output lines;
 - a plurality of packet switched output modules, said packet switched output modules together having a plurality of inputs, each of said packet switched output module inputs connected to respective output port of said second portion of said single stage switching fabric, and together having a plurality of outputs, each of said packet switched output module outputs connected to a respective packet switched output line of said plurality of packet switched output lines;
- means for hunting a path through said switch fabric to a desired circuit switched output line for communication on each circuit switched input line; and
- means for hunting a path through said switch fabric to a desired packet switched output line for a packet on each packet switched input line.
2. The switch as set forth in claim 1, wherein:
 - said circuit switched communication path hunting means includes a first out of band, controller; and
 - said packet switched communication path hunting means includes a second out of band, controller.
 3. The switch as set forth in claim 2, wherein said switch fabric is partitioned into multiple pipes.

24

4. The switch as set forth in claim 3, wherein said out-of-band controllers roll requests which have been denied a path through a first pipe to a second pipe.

5. The switch as set forth in claim 3, wherein said out-of-band controllers roll requests which have been denied a path through a first pipe and a second pipe to a third pipe.

6. The switch as set forth in claim 3, wherein said out-of-band controllers roll requests which have been denied a path through a first pipe, a second pipe and a third pipe to a fourth pipe.

7. The switch as set forth in claim 3, wherein said second out-of-band controller assigns an order of preference to communication packets.

8. The switch as set forth in claim 1, wherein a size of said first portion of said switch fabric is in a range of zero to one hundred per cent.

9. The switch as set forth in claim 8, wherein a length of each communication packet may vary from one packet to another.

10. The switch as set forth in claim 1, wherein a length of each communication packet may vary from one packet to another.

11. A packet switch for switching a telecommunication packet from a plurality of input lines to a plurality of output lines, comprising:

a plurality of input interfaces, each having an input port connected to a respective input line of said plurality of input lines, and each of said input interfaces having an output port;

a network for switching a plurality of I input ports to a plurality of P output ports;

each of said plurality of input interface output ports is fanned out to a respective group of F of said I input ports of said network;

said network having a plurality of C pipes, where C is an integer of a value equal to P/I ;

a plurality of output modules, said output modules together having a plurality of inputs, each of said output module inputs connected to respective output port of said plurality of P output ports, and together having a plurality of outputs, each of said output module outputs connected to a respective output line of said plurality of output lines;

each pipe of said C pipes having a path from each of the plurality of inputs lines that is connectable to a respective output line of the plurality of output lines;

a spare pipe for on-line replacement of any one of said C pipes which is faulty; and

means for hunting a path through said packet switch for a telecommunication packet.

* * * * *

DISTRIBUTION LIST

addresses	number of copies
AIR FORCE RESEARCH LABORATORY/SNDR JAMES L DAVIS 25 ELECTRONIC PKY ROME NY 13441-4515	5
LUCENT TECHNOLOGIES 2000 NORTH NAPERVILLE RD NAPERVILLE IL 60566	5
AFRL/IFOIL TECHNICAL LIBRARY 26 ELECTRONIC PKY ROME NY 13441-4514	1
ATTENTION: DTIC-OCC DEFENSE TECHNICAL INFO CENTER 8725 JOHN J. KINGMAN ROAD, STE 0944 FT. BELVOIR, VA 22060-6218	2
ADVANCED RESEARCH PROJECTS AGENCY 3701 NORTH FAIRFAX DRIVE ARLINGTON VA 22203-1714	1
RELIABILITY ANALYSIS CENTER 201 MILL ST. ROME NY 13440-8200	1
ROME LABORATORY/C3AB 525 BROOKS RD ROME NY 13441-4505	1
ATTN: GWEN NGUYEN GIDEP P.O. BOX 8000 CORONA CA 91718-8000	1

AFIT ACADEMIC LIBRARY/LDEE 1
2950 P STREET
AREA B, BLDG 642
WRIGHT-PATTERSON AFB OH 45433-7765

ATTN: R.L. DENISON 1
WRIGHT LABORATORY/MLPD, BLDG. 651
3005 P STREET, STE 6
WRIGHT-PATTERSON AFB OH 45433-7707

WRIGHT LABORATORY/MTM, BLDG 653 1
2977 P STREET, STE 6
WRIGHT-PATTERSON AFB OH 45433-7739

ATTN: GILBERT G. KUPERMAN 1
AL/CFHI, BLDG. 248
2255 H STREET
WRIGHT-PATTERSON AFB OH 45433-7022

ATTN: TECHNICAL DOCUMENTS CENTER 1
DL AL HSC/HRG
2698 G STREET
WRIGHT-PATTERSON AFB OH 45433-7604

AIR UNIVERSITY LIBRARY (AUL/LSAD) 1
600 CHENNAULT CIRCLE
MAXWELL AFB AL 36112-6424

US ARMY SSDC 1
P.O. BOX 1500
ATTN: CSSD-IM-PA
HUNTSVILLE AL 35807-3801

TECHNICAL LIBRARY D0274(PL-TS) 1
SPAWARSYSCEN
53560 HULL STREET
SAN DIEGO CA 92152-5001

NAVAL AIR WARFARE CENTER 1
WEAPONS DIVISION
CODE 4BL000D
1 ADMINISTRATION CIRCLE
CHINA LAKE CA 93555-6100

SPACE & NAVAL WARFARE SYSTEMS CMD 2
ATTN: PMW163-1 (R. SKIANO) RM 1044A
53560 HULL ST.
SAN DIEGO, CA 92152-5002

SPACE & NAVAL WARFARE SYSTEMS 1
COMMAND, EXECUTIVE DIRECTOR (PD13A)
ATTN: MR. CARL ANDRIANI
2451 CRYSTAL DRIVE
ARLINGTON VA 22245-5200

COMMANDER, SPACE & NAVAL WARFARE 1
SYSTEMS COMMAND (CODE 32)
2451 CRYSTAL DRIVE
ARLINGTON VA 22245-5200

CDR, US ARMY MISSILE COMMAND 2
REDSTONE SCIENTIFIC INFORMATION CTR
ATTN: AMSMI-RD-CS-R, DOCS
REDSTONE ARSENAL AL 35898-5241

ADVISORY GROUP ON ELECTRON DEVICES 1
SUITE 500
1745 JEFFERSON DAVIS HIGHWAY
ARLINGTON VA 22202

REPORT COLLECTION, CIC-14 1
MS P364
LOS ALAMOS NATIONAL LABORATORY
LOS ALAMOS NM 87545

AEDC LIBRARY 1
TECHNICAL REPORTS FILE
100 KINDEL DRIVE, SUITE C211
ARNOLD AFB TN 37389-3211

COMMANDER 1
USAISC
ASHC-IMD-L, BLDG 61801
FT HUACHUCA AZ 85613-5000

US DEPT OF TRANSPORTATION LIBRARY 1
FB10A, M-457, RM 930
800 INDEPENDENCE AVE, SW
WASH DC 22591

AWS TECHNICAL LIBRARY
859 BUCHANAN STREET, RM. 427
SCOTT AFB IL 62225-5118

1

AFIWC/MSY
102 HALL BLVD, STE 315
SAN ANTONIO TX 78243-7016

1

SOFTWARE ENGINEERING INSTITUTE
CARNEGIE MELLON UNIVERSITY
4500 FIFTH AVENUE
PITTSBURGH PA 15213

1

NSA/CSS
K1
FT MEADE MD 20755-6000

1

ATTN: DM CHAUHAN
DCMC WICHITA
271 WEST THIRD STREET NORTH
SUITE 6000
WICHITA KS 67202-1212

1

PHILLIPS LABORATORY
PL/TL (LIBRARY)
5 WRIGHT STREET
HANSCOM AFB MA 01731-3004

1

ATTN: EILEEN LADUKE/D460
MITRE CORPORATION
202 BURLINGTON RD
BEDFORD MA 01730

1

OUSDC(P)/DTSA/DUTD
ATTN: PATRICK G. SULLIVAN, JR.
400 ARMY NAVY DRIVE
SUITE 300
ARLINGTON VA 22202

2

RICHARD PAYNE
AIR FORCE RESEARCH LAB/SNH
HANSCOM AFB, MA 01731-5000

1

JOSEPH P. LORENZO, JR.
AIR FORCE RESEARCH LAB/SNHC
HANSCOM AFB, MA 01731-5000

1

JOSEPH L. HORNER
AIR FORCE RESEARCH LAB/SNHC
HANSCOM AFB, MA 01731-5000

1

RICHARD A. SOREF
AIR FORCE RESEARCH LAB/SNHC
HANSCOM AFB, MA 01731-5000

1

JOHN J. LARKIN
AIR FORCE RESEARCH LAB/SNHX
HANSCOM AFB, MA 01731-5000

1

ALBERT A. JAMBERDINO
AIR FORCE RESEARCH LAB/IFED
32 HANGAR RD
ROME NY 13441-4114

1

AIR FORCE RESEARCH LAB/SND
25 ELECTRONIC PKY
ROME NY 13441-4515

1

JOANNE L. ROSSI
AIR FORCE RESEARCH LAB/SNW
25 ELECTRONIC PKY
ROME NY 13441-4515

1

NY PHOTONIC DEVELOPMENT CORP
MVCC ROME CAMPUS
UPPER FLOYD AVE
ROME, NY 13440

1

ROBERT T. KEMERLEY
AIR FORCE RESEARCH LABORATORY/SND
2241 AVIONICS CIRCLE, RM C2G69
WRIGHT-PATTERSON AFB OH 45433-7322

1